

Massachusetts Institute of Technology

Artificial Intelligence Laboratory

AI Memo 449

January, 1978

THE GENETIC EPISTEMOLOGY OF RULE SYSTEMS

Ira P. Goldstein

I shall describe a model of the evolution of rule-structured knowledge that serves as a cornerstone of our development of computer-based coaches. The key idea is a graph structure whose nodes represent rules, and whose links represent various evolutionary relationships such as generalization, correction, and refinement. This graph guides both student modelling and tutoring as follows: the coach models the student in terms of nodes in this graph, and selects tutoring strategies for a given rule on the basis of its genetic links. It also suggests a framework for a theory of learning in which the graph serves as a memory structure constructed by the student by means of processes corresponding to the various links. Given this framework, a learning complexity measure can be defined in terms of the topology of the graph.

Keywords: Information Processing Psychology, Learning, Knowledge representation, Computer Aided Instruction, Artificial Intelligence, Cognitive Science.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. It was supported by the National Science Foundation under grant SED77-19279.

Accepted for publication by the International Journal of Man-Machine Studies.

CONTENTS

Acknowledgments	iv
1. Introduction	5
1.1 A Learner-based Paradigm for AICAI is evolving.	5
1.2 A graph representation of the syllabus has roots in AICAI research.	6
2. Wumpus serves as an experimental domain.	9
2.1 Definition of the Wumpus Game	9
2.2 Wumpus AICAI Tutors	10
3. The genetic graph formalizes the syllabus.	12
3.1 Genetic links specify evolutionary relationships between rules.	12
3.2 Genetic graphs are being explored in a student simulation testbed.	15
4. Extensions to the genetic graph.	17
4.1 The extended GG groups related knowledge into islands.	17
4.2 The extended GG represents the justifications of rules.	19
4.3 The extended GG represents planning knowledge.	19
4.4 The extended GG is a network of islands.	22
5. The genetic graph is a basis for tutoring.	23
5.1 The genetic graph suggests the tutoring topic.	23
5.2 The genetic graph supplies multiple explanations.	25
5.3 Tutoring using an Extended Syllabus Representation	26
5.4 The genetic graph does not solve the tutoring problem.	27
6. The genetic graph is a basis for modelling.	27
6.1 The student knowledge model overlays the nodes of the genetic graph.	28
6.2 The genetic graph guides the construction of the model.	29
6.3 The student learning model overlays the links of the genetic graph.	31
6.4 The genetic graph does not solve the modelling problem.	31

7. The genetic graph is a basis for learning.	33
7.1 The student is an active agent.	35
7.2 A genetic graph for learning skills is possible.	35
7.3 A belief measure can be defined on the genetic graph.	37
7.4 The genetic graph topology provides a learning complexity measure.	39
7.5 Designing simulated students is a research methodology.	40
7.6 The genetic graph is not a complete theory of learning.	41
8. Conclusions	41
9. References	43

Acknowledgments

This paper has evolved from many fruitful conversations with members of the Cognitive Computing Group at the MIT AI Lab (in particular Greg Clemenson, Mark Miller, Sandy Schoichet, Bob Sjoberg, Bill Swartout, Barbara White, Kurt van Lehn, Bruce Roberts, Jim Stansfield, and Steve Rosenberg) and with members of the ICAI group at Bolt, Beranek and Newman (in particular John Seely Brown and Dick Burton).

1. Introduction

1.1 A Learner-based Paradigm for AICAI is evolving.

The 1970's has seen the evolution of a new generation of computer-aided instructional programs based on the inclusion of AI-based expertise within the CAI system. These systems surmount the restrictive nature of older script-based CAI by supplying "reactive" learning environments which can analyze a wide range of student responses by means of an embedded domain-expert. Examples are AICAI tutors for geography [Car70], electronics [Bro73], set theory [Smi75], Nuclear Magnetic Resonance spectroscopy [Sle75], and mathematical games [Bur76, Gol77a].

However, while the inclusion of domain expertise is an advance over earlier script-based CAI, the tutoring theory embedded within these benchmark programs for conveying this expertise is elementary. In particular, they approach teaching from a subset viewpoint: expertise consists of a set of facts or rules. The student's knowledge is modelled as a subset of this knowledge. Tutoring consists of encouraging the growth of this subset, generally by intervening in situations where a missing fact or rule is the critical ingredient needed to reach the correct answer.

This is, of course, a simplification of the teaching process. It has allowed research to focus on the critical task of representing expertise. But the subset viewpoint fails to represent the fashion in which new knowledge evolves from old by such processes as analogy, generalization, debugging, and refinement.

This paper explores the genetic graph as a framework for representing procedural knowledge from an evolutionary viewpoint,¹ thereby contributing to the movement of AICAI from an

1. A potential confusion in terminology may occur here. The term "genetic" is often equated with heredity. However, I use it here in its older sense, namely, the genetic method is the study of the origins and development of a phenomena. This paper is an exercise in Genetic Epistemology, the study of the origin and development of knowledge. This enterprise has been articulately advocated by Piaget [Pia70], who considers it the foundation on which psychology should be based.

expert-based to a *learner-based* paradigm.² After introducing our experimental domain, the mathematical game Wumpus, and describing an expert-based coach which we have implemented, I define the genetic graph and describe how it can improve the range of tutoring advice that the AICAI system can provide and the accuracy of the model that the system builds of the learner. Fig. 1 illustrates its central role. I then discuss in greater detail the model of the learner implicit in the genetic graph. By articulating this model, I am able to suggest a measure of *learning complexity* in terms of the topological properties of the graph. I conclude with a suggestion for reformulating traditional Piagetian notions of accommodation, assimilation and equilibration in terms of our procedural epistemology.

While I shall describe a student simulation testbed which we have implemented to test various genetic graph formulations, this paper is largely exploratory. Its purpose is to serve as a critique of existing expert-based AICAI systems, in particular our Wumpus coach, and a proposal for an improved "learner-based" design.

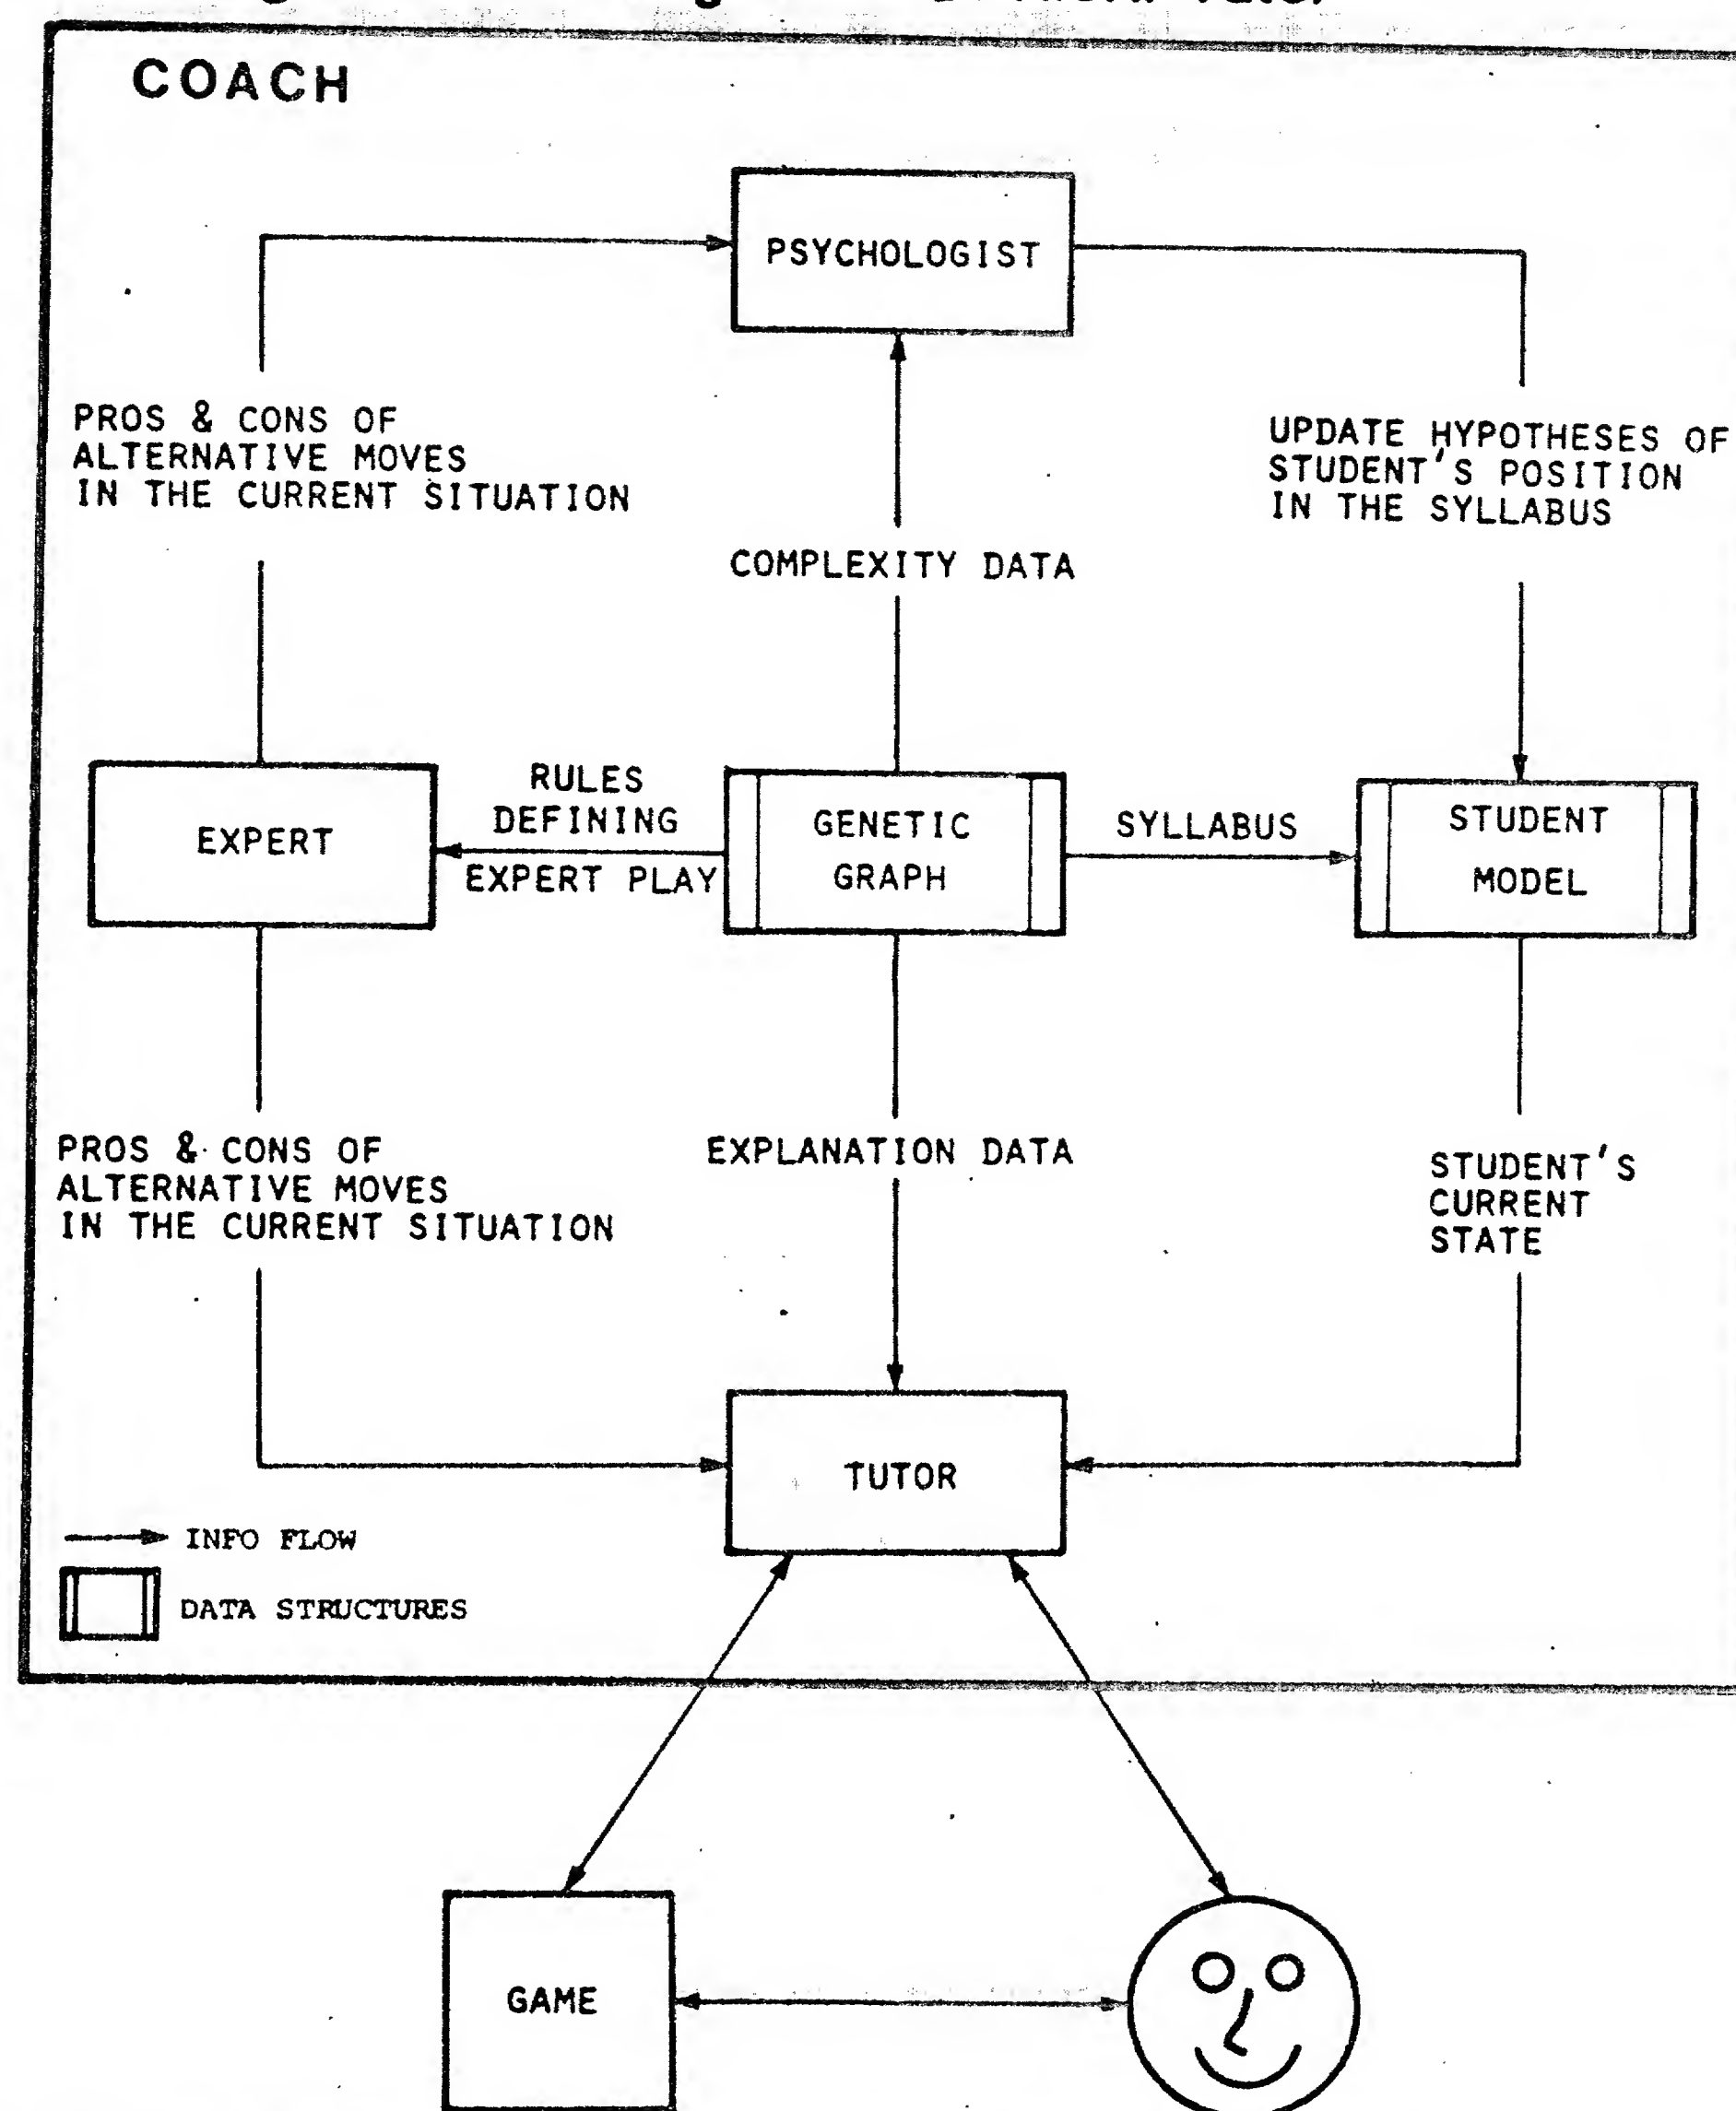
1.2 A graph representation of the syllabus has roots in AICAI research.

Scholar [Car70], the earliest of the AICAI tutors, employed a graph (semantic net) representation for declarative facts about geography. The graph, however, encoded only domain specific relationships; it did not embody a series of progressively more refined levels of geography knowledge linked by various evolutionary relationships.³

2. There are other dimensions to this paradigm shift that include: (1) more sophisticated modelling of the student's knowledge and learning style [Bur76, Bro77a, Carr77b], (2) widening the communication channel from student to teacher via natural language interfaces [Bur77], and (3) developing a theory of teaching skills [Col75].

3. Scholar might be extended in this fashion, especially if employed with younger children whose theory of the world may not already be stabilized in the expert form embodied by Scholar.

Fig. 1. Block Diagram of an AICAI Tutor



SOPHIE-1 [Bro73], the next major AICAI milestone, was an expert-based system for the more complex domain of electronic troubleshooting. SOPHIE-1 compared a student's troubleshooting hypotheses for an electronic circuit with that of its embedded expert and offered advice when the student's analysis went astray. It employed a procedural rather than a network representation for its electronics knowledge, but this representation was largely a black box. SOPHIE-1 did not have access to a detailed, modular, human-oriented representation of troubleshooting skills. Nor did it have a representation for the genesis of these skills.

SOPHIE-2, now under development, will incorporate a modular, anthropomorphic representation for the expert's knowledge [Dek76]. This structured expertise serves as a better foundation for expert-based tutoring, but still is not a model of how the student evolved to that level of competence.

BUGGY [Bro77a], a program for building procedural models of a student's arithmetic skills, does incorporate both a graph representation for the basic skills and some evolutionary relationships. The basic skill representation is a graph with links representing the skill/subskill relationships. The evolutionary component consists of "deviation" links to "buggy" versions of the various skills.

BIP-II [Wes77], a tutor for programming skills, again employs a network for the basic skill representation, but embodies a different set of evolutionary relationships. There are links for representing analogy, generalization, specialization, prerequisite, and relative difficulty relations. The BIP-II skill network, however, does not include deviation links nor define an operational expert for the programming domain. Rather it employs author-supplied exercises attached to the relevant skills in the network.⁴

The genetic graph is a descendant of these network representations. Its nodes are the procedural skills of players of varying proficiency and its links include the analogy, specialization, generalization and prerequisite relations of BIP-II and the deviation relationships of BUGGY.⁵

4. MALT [Kof75], a tutor for machine language programming, does include an "expert" for problems composed from a limited set of skills and solved in a tutor-prescribed order. However, MALT's syllabus of skills are related only by the probability with which MALT includes them in a system-generated problem, and not by any evolutionary links. Hence, MALT does not have BIP's ability to choose a problem based on its evolutionary relationship to the student's current knowledge state.

5. The skill nodes themselves, corresponding to rules of the form: "if C1 & C2 & ..., do A1 & A2 & ...", could be expanded into more primitive networks of conditions and conjunction nodes similar to those employed in BUGGY and BIP-II, but I do not discuss that extension in this paper. Instead, I concentrate on describing the evolutionary relationships between skills.

2. Wumpus serves as an experimental domain.

Designing coaches for the maze exploration game Wumpus [Yob75] has proven to be a profitable experimental domain because the game exercises basic skills in logic and probability.⁶ This section defines our version of the game⁷ and describes two expert-based coaches which have been previously implemented for it. The next section formulates an evolutionary epistemology of the knowledge required for skilled play, providing the basis for an improved "learner-based" design.

2.1 Definition of the Wumpus Game

The player is initially placed somewhere in a warren of caves with the goal of slaying the Wumpus. The difficulty in finding the beast arises from the existence of dangers in the warren -- bats, pits and the Wumpus itself. Pits and the Wumpus are fatal; bats move the player to a random cave elsewhere in the warren. But the player can infer the probable location of dangers from warnings he receives. The Wumpus can be sensed two caves away, pits and bats one cave away. Victory results from shooting an arrow into the Wumpus's lair; defeat if the arrows are fruitlessly exhausted.

Becoming skilled poses a nontrivial learning experience for most children and adults:⁸ locating multiple dangers in a randomly connected warren of twenty or more caves can be complex.

6. Our group is also exploring evolutionary epistemologies for other domains ranging from elementary programming to airplane flying.

7. Yob's original game was played on the graph of a dodecahedron. Our version is a generalization involving a variable maze geometry, a variable number of dangers, and a variable warning distance for each danger.

8. By nontrivial, I mean that the experience is on the order of hours rather than minutes and that some players (given no coaching) fail to acquire certain skills after many hours of play. This is based on informal observations of over thirty players ranging in age from grade school children to adults.

Hence, the game provides a useful problem domain for developing a theory of the evolution of procedural skills.

2.2 Wumpus AICAI Tutors

In 1976 we developed WUSOR-I [Sta76], an expert-based coach. Skilled play was analyzed in terms of rules such as these:

Positive Evidence: A warning implies that a danger is in a neighboring cave.

Elimination: If a cave has a warning and all but one of its neighbors are known to be safe, then the danger is in the remaining neighbor.

Multiple Evidence: Multiple warnings increase the likelihood that a given cave contains a danger.

Below is an interaction with the WUSOR-I coach which illustrates the fashion in which the coach explains a rule if its employment would result in a better move than the one chosen by the student. The student's responses are indented and underlined.

You are at cave 15 with neighbors 4, 14 and 0. Brrr! There is a draft. You are near a pit. What now?

1 > 4


You are at cave 4 with neighbors 15, 14 and 2. Brrr! There is a draft. Squeak! A bat is near. (The game state is show in fig. 2.) What now?

2 > 14

(E1) Mary, it isn't necessary to take such large risks with pits. There is multiple evidence of a pit in cave 14 which makes it quite likely that cave 14 contains a pit. It is less likely that cave 0 contains a pit. Hence, we might want to explore cave 0 instead. Do you want to take back your move?

WUSOR-I was insensitive to the relative difficulty of the various Wumpus skills. In 1977 we took our first steps toward an evolutionary epistemology with WUSOR-II [Carr77a], wherein the rule set was divided into five phases representing increasing skill at the task.

Fig. 2. The Wumpus game state after two moves



Phase 1: Rules for visited, unvisited and fringe caves.

Phase 2: Rules for possibly dangerous, definitely dangerous and safe caves.

Phase 3: Rules for single versus multiply dangerous caves.

Phase 4: Rules for "possibility sets", i.e. keeping track of the sources of dangers.

Phase 5: Rules for numerical evidence.

The tutor did not describe the rule of a particular level of play until it believed the student was familiar with the rules of the preceding levels.⁹

These phases constituted a coarse genetic epistemology, better than the completely unordered approach of WUSOR-I, but still far from a detailed platform on which to build new knowledge from old in the student's mind. WUSOR-III, now being implemented, addresses this limitation. It has evolved from WUSOR-II by defining a set of symbolic links between rules that characterize such relationships as analogy, refinement, correction, and generalization.¹⁰ The result is that the "syllabus" of the coach has evolved from an unordered skill set to a *genetic graph* of skills linked by their evolutionary relationships.

9. [Carr77b] describes the mechanisms by which it estimated the student's position in the syllabus.

10. It was also necessary to increase the grain of the rules. WUSOR-II rules were too coarse, and hence obscured certain evolutionary relationships.

3. The genetic graph formalizes the syllabus.

The "genetic graph" (GG) formalizes the evolution of procedural rules by representing the rules as nodes and their interrelationships as links. In this section I discuss four of these relationships -- generalization/specialization, analogy, deviation/correction, and simplification/refinement -- and provide examples of their occurrence in the Wumpus syllabus. I also describe a student simulation testbed which we have implemented to explore the consequences of different rule formulations. In the next section, I consider what kinds of knowledge are not properly represented by a graph of rules, and propose appropriate extensions.

3.1 Genetic links specify evolutionary relationships between rules.

R' is a *generalization* of R if R' is obtained from R by quantifying over some constant.¹¹ *Specialization* is the inverse relation. In the Wumpus syllabus, for each trio of specialized rules for bats, pits and the Wumpus, there is usually a common generalization in terms of warnings and dangers.¹² Fig. 3 illustrates such a cluster for rule 2.2 which represents the deduction: "a warning implies that the neighbors of the current cave are dangerous."

R' is *analogous* to R if there exists a mapping from the constants of R' to the constants of R . This is the structural definition employed by Moore and Newell [Moor73]. Of course, not all analogies defined in this fashion are profitable. However, the GG is employed to represent those that are.

11. This is a standard predicate calculus definition, applied here to quantifying over formulas representing rules rather than logical statements.

12. In one version of Wumpus, the wumpus warning propagates only one cave. In this case, bats, pits and the wumpus are exactly analogous. In more complex versions, the Wumpus is no longer exactly analogous. Hence, the analogies to bats and pit rules are in fact restricted cases or outright deviations. We represent this in the GG explicitly, thereby giving the coach an expectation for the traps the student will encounter.

Fig. 3. A Region of the Genetic Graph

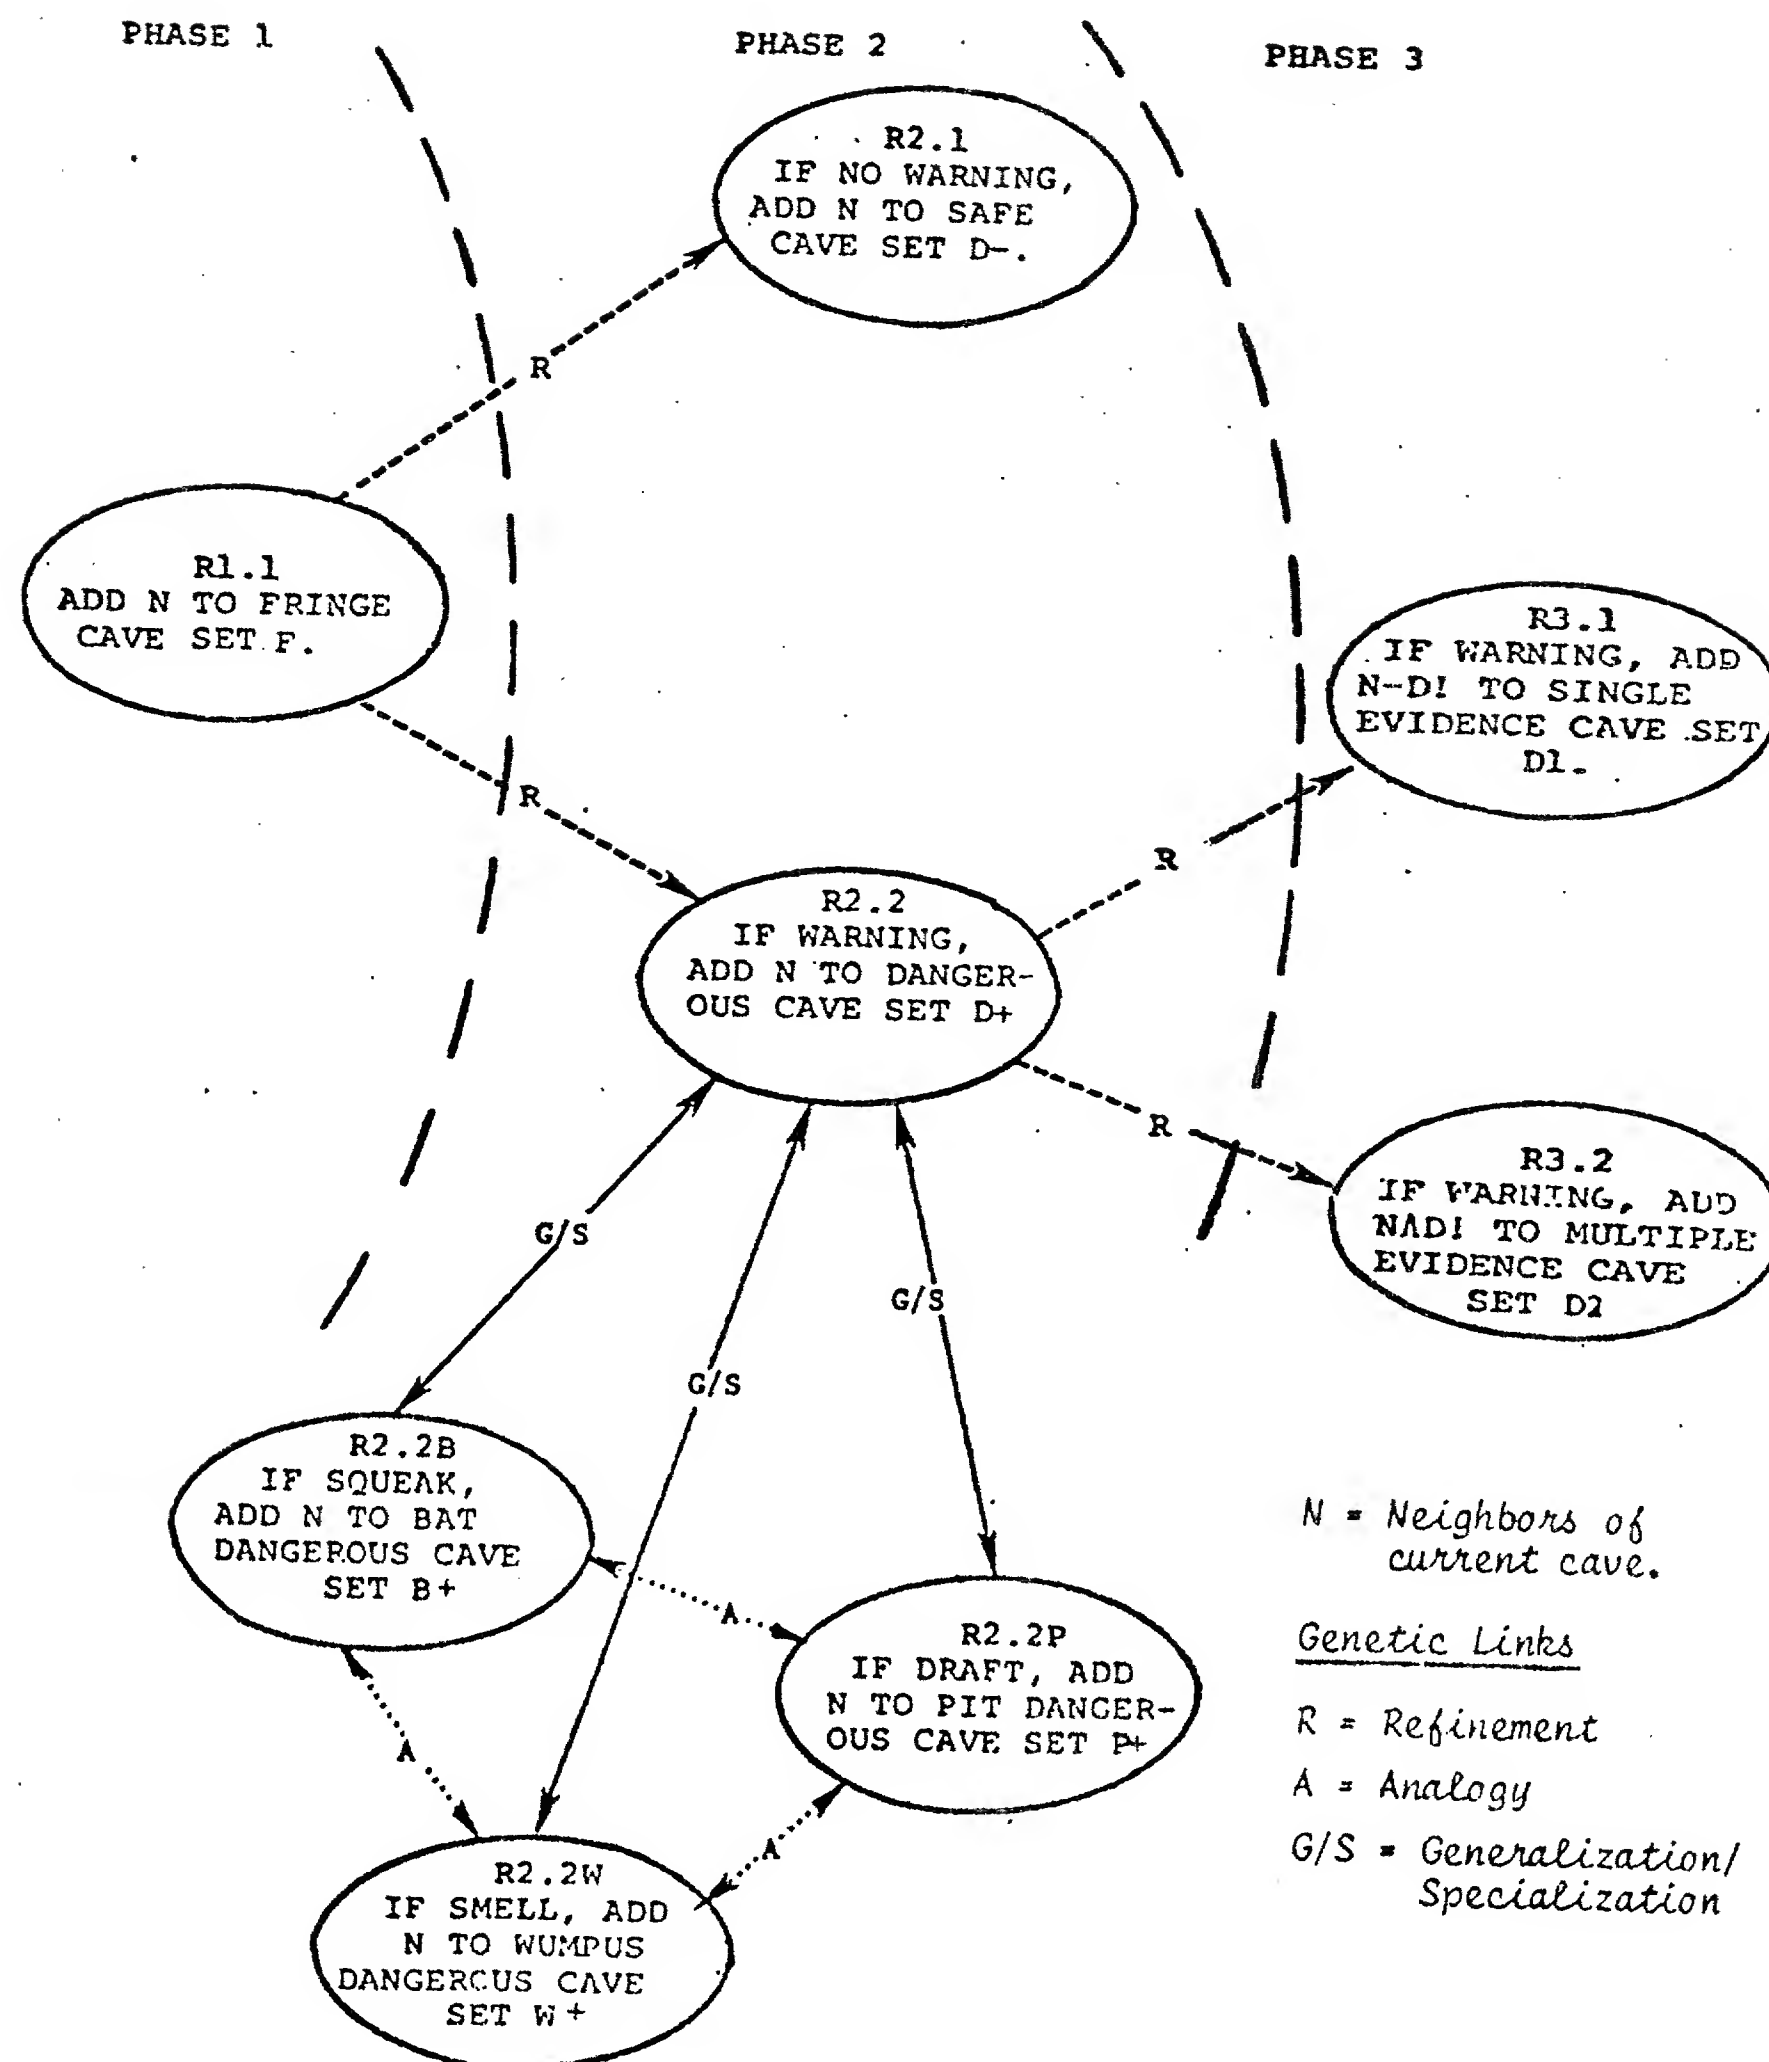


Fig. 3 illustrates analogy links between the specialization trio of R2.2. For example, mapping SQUEAK to DRAFT and B+ (the set of caves risking BATS) to P+ (the set of caves risking PITS) defines the analogy mapping between R2.2B and R2.2P. The similar nature of dangers makes clusters of this kind (one generalization and three specializations all connected by analogy links) common in the Wumpus world. As we shall discuss in sec. 5, identifying such densely linked clusters provides teaching leverage by providing multiple methods of explanation (one per link) for each constituent rule.

R' is a *refinement* of R if R' manipulates a subset of the data manipulated by R on the basis of some specialized properties. *Simplification* is the inverse relation. This relation represents the evolution of a rule to take account of a finer set of distinctions. The Wumpus syllabus contains five major refinements corresponding to the five Wusor-II phases. Fig. 3 illustrates the refinement of the rule $R1.1$ through phases 1, 2 and 3. $R2.1$ and $R2.2$, for example, refine $R1$'s treatment of the fringe caves by distinguishing between and safe and dangerous subsets. $R3.1$ and $R3.2$, in turn, refine the dangerous subset into single and multiply dangerous categories.

R' is a *deviation* of R if R' has the same purpose as R but fails to fulfill that purpose in some circumstances. *Correction* is the inverse relation. Deviations arise naturally in learning as the result of simplifications, overgeneralizations, mistaken analogies, and so on. While any rule can have deviant forms, the GG is used to record the more common errors.¹³

A deviant Wumpus rule is: "If there is multiple evidence that a cave contains a pit, then that cave definitely contains a pit." The debugged rule includes the additional condition that there is only one pit in the warren. The deviation has a natural genetic origin: it is a reasonable rule in the early stages of Wumpus play when the game is simplified by the coach to contain only one of each danger.

13. The deviant skills recorded in the GG account for errors arising from the correct application of incorrect rules. There is another class of errors arising from the incorrect application of correct rules. These are errors arising from such causes as the occasional failure to check all preconditions of a rule, the misreading of data, or confusion in the bookkeeping associated with a search process. Sleeman [Sle77] explores some errors of this class in his construction of a coach which analyzes a student's description of his algorithm. Sleeman's coach, however, does not have a representation for deviant or simplified versions of the algorithm to be tutored: indeed, he assumes that the student is familiar with the basic algorithm. A possible extension of his system would be to include a GG representing evolutionary predecessors of the skilled expert.

3.2 Genetic graphs are being explored in a student simulation testbed.

The Wumpus GG currently contains about 100 rules and 300 links.¹⁴ We are currently testing the reasonableness of this graph by means of a "Student Simulation Testbed".^{15,16} In this testbed, the performance of various simulated students, defined in terms of different regions of the GG, is being examined. These students correspond to different evolutionary states. Fig. 4 is the comparative trace of two students corresponding to the mastery of phases 2 and 3 respectively.

The "WHY" messages of fig. 4 are printed by the student simulator as the rules defining a student are executed. The comments inside cave boxes represent hypotheses of the simulated student regarding that cave. The balloons reflect the differing hypotheses of the two students regarding bat evidence for caves J and H.

The phase 2 student (dotted path) does not know the multiple evidence heuristic. Hence, he does not realize that cave J is to be preferred over cave H. While he understands that they both risk bats, he makes no further distinction. Thus, he randomly selects from these two possibilities, unfortunately choosing the riskier H. The phase 3 student (dashed path) recognizes multiple (BAT2) evidence as more risky than single (BAT1) evidence and therefore selects the safer cave J.

14. These statistics are based on an explicit representation of each generalization, its specializations and their common deviations. It is possible for the graph to be less extensive if procedures for generating common deviations and specializations are supplied. This is the approach we shall eventually employ. Specializations are simple to generate. Deviations are suggested by the common bug types enumerated by such work as my own analysis of Logo programs [Gol75], Sussman's analysis of Blocks world programs [Sus75], and Stevens and Collins's study of bugs in causal reasoning [Ste77]; or they can be induced, for simple cases, by analyzing the student's performance [Sel74, Gol75, Bro77b]. However, my current research strategy has been to make the graph explicit, in order to understand its form. The next stage will include the extension to expanding the graph dynamically.

15. The testbed serves other purposes as well. Simulated students can be used to test the modelling and tutoring of teaching systems [Carr77a, Sel77, Wes77]. They can also serve as models of real students, and hence can yield insight for a human teacher observing their performance [Bro77a, Gol77b].

16. Following this testing period, WUSOR-II will be converted to incorporate the GG. The expected improvement in modelling and tutoring is the subject of sections 5 and 6.

Fig. 4. Divergent Behavior of Two Simulated Students




Fig. 4 is a composite of the graphic output for the two students. The testbed only executes a single student at a time. It does not generate balloons nor place the "WHY" messages on the Warren itself.

Expert-based CAI allows only for the definition of "simulated students" formed from subsets of the expert's skills. The power of the GG to broaden the tutor's understanding of the task is evident from the testbed: the GG permits not only the creation of subset students, but also students formed from specializations, deviations, and simplifications of the expert's rules.

Nevertheless, it must be stressed that the evolutionary relations discussed here remain both underspecified and incomplete. There are many kinds of analogies, generalizations, and corrections. There are also other kinds of evolutionary processes for acquiring knowledge: learning by being told, learning by induction from past examples, and learning by deduction from old rules. The

next chapter explores one of the directions in which the GG must be extended to be an adequate representation for the evolution of a student's knowledge.

4. Extensions to the genetic graph.

The preceding chapter defined a set of genetic relationships between individual rules. In this chapter, we extend the genetic graph to incorporate genetic relationships between groups of rules and between rules and the declarative facts that explain and justify their behavior.

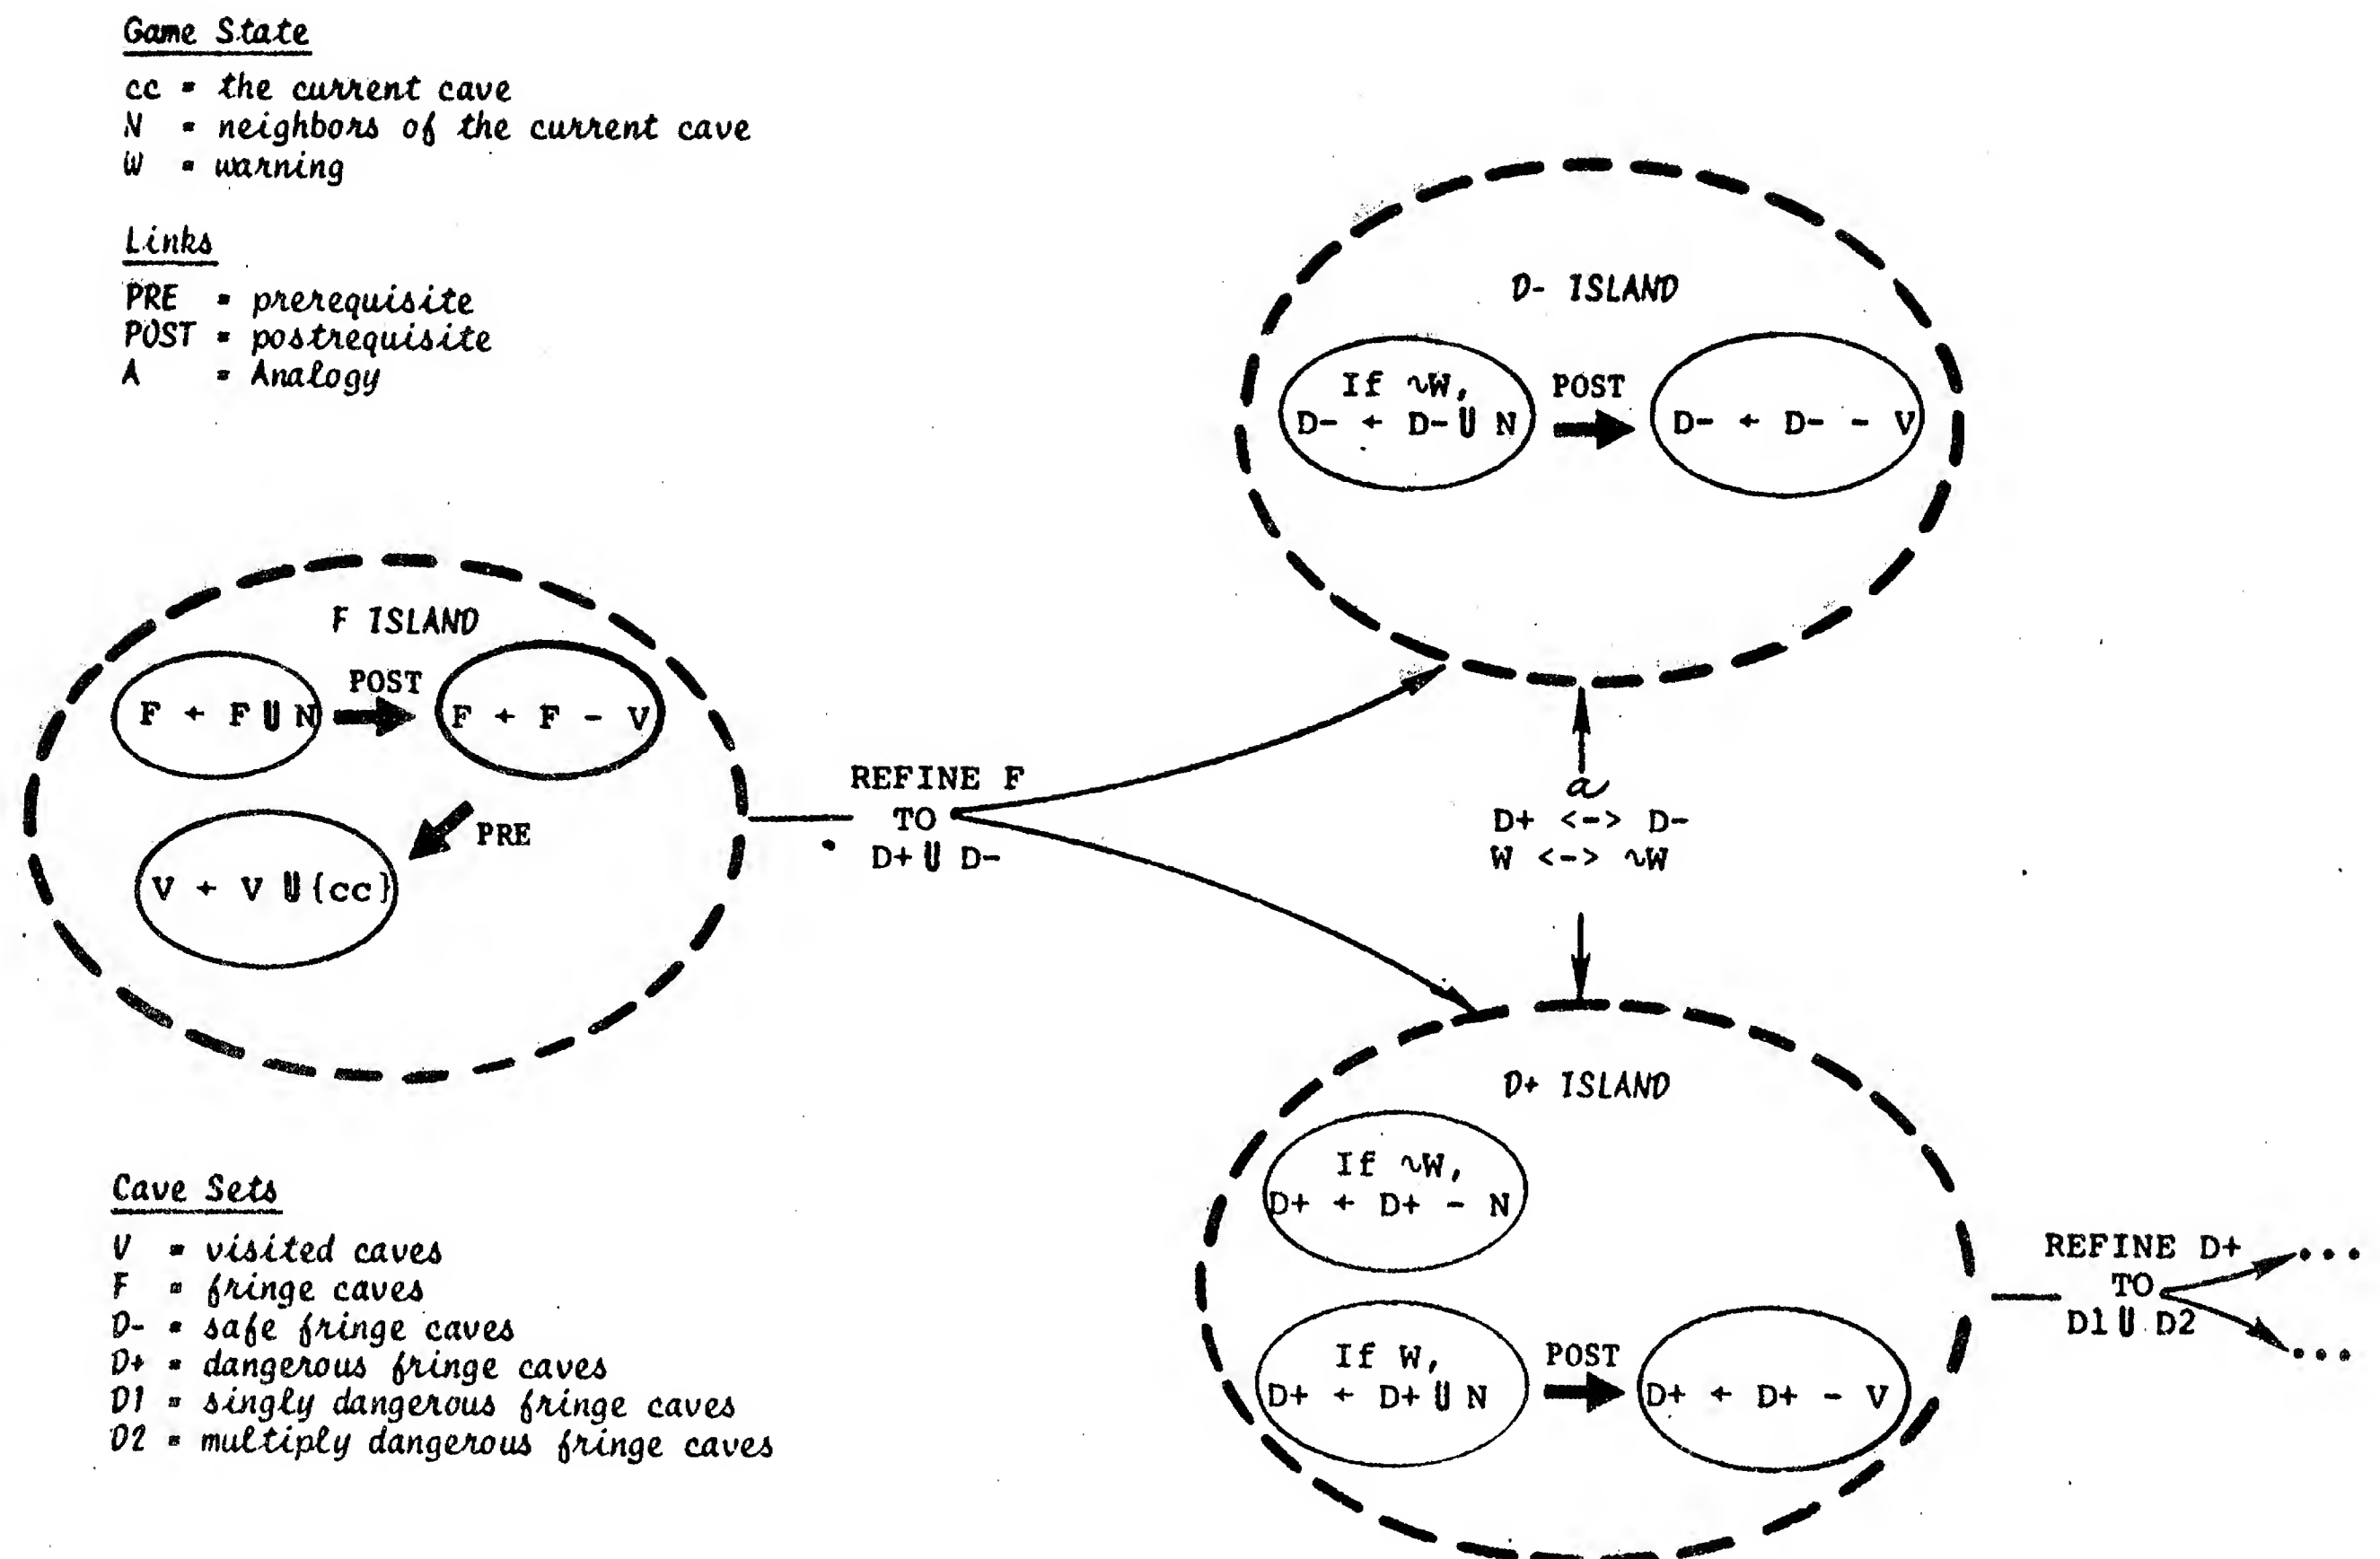
4.1 The extended GG groups related knowledge into islands.

Our first order theory of a GG has the limitation that rules which are closely related and generally learned as a group are not so represented. To address this limitation, rules are grouped into *islands*. A natural criterion for forming islands is to group rules that have the same goal. For Wumpus, this translates into grouping rules which manipulate the same kind of evidence. This is illustrated in fig. 5.

For example, the D^+ island contains the rules which manipulate D^+ , the set of possibly dangerous caves. One rule subtracts the neighbors of the current cave from D^+ if there is no warning. The complementary rule is also present: it adds to D^+ the neighbor set if there is a warning. The third rule in the D^+ island subtracts the visited caves from D^+ . It insures that D^+ contains only fringe caves. (A prerequisite planning link exists between these last two rules which is explained later in this section.)

Islands allow the coach to tutor the student in terms of an overall concept for a group of rules and to model the student in terms of his possession of the conceptual base underlying a rule set. Just as an analogy between two rules can be explained, so too can an analogy between two islands of rules. Fig. 5 illustrates this with an analogy link between the safe and dangerous islands.

Fig. 5. The Extended Genetic Graph is a Network of Islands



The acquisition of a group of skills is a natural learning episode since acquiring the island is a local task -- the rules all follow from a single concept. But moving to the next island requires a new conceptual base. To explore this movement, the simulated student testbed allows macro instructions which add entire islands of rules to the simulated student being constructed.

- > (Define_student 5 (island F) (island D-) (island D+))
 Student 5 defined. ;This is the phase 2 player of fig. 4.
- > (Define_student 6 (student 5) (island D1) (island D2))
 Student 6 defined. ;This is the phase 3 player of fig. 4.

4.2 The extended GG represents the justifications of rules.

The GG as a representation of knowledge is still incomplete. Rules by themselves do not describe the declarative knowledge that explains and justifies their behavior. For Wumpus, this declarative knowledge includes the definition of the evidence sets and axiomatic statements of their properties. Fig. 6 shows the declarative facts listed below linked to the various groups of rules whose behavior they justify.

The fringe F is the union of $D+$ and $D-$ (the safe caves).

A warning implies that some of the neighbors of the current cave are members of $D+$.

$D-$ and $D+$ are disjoint.

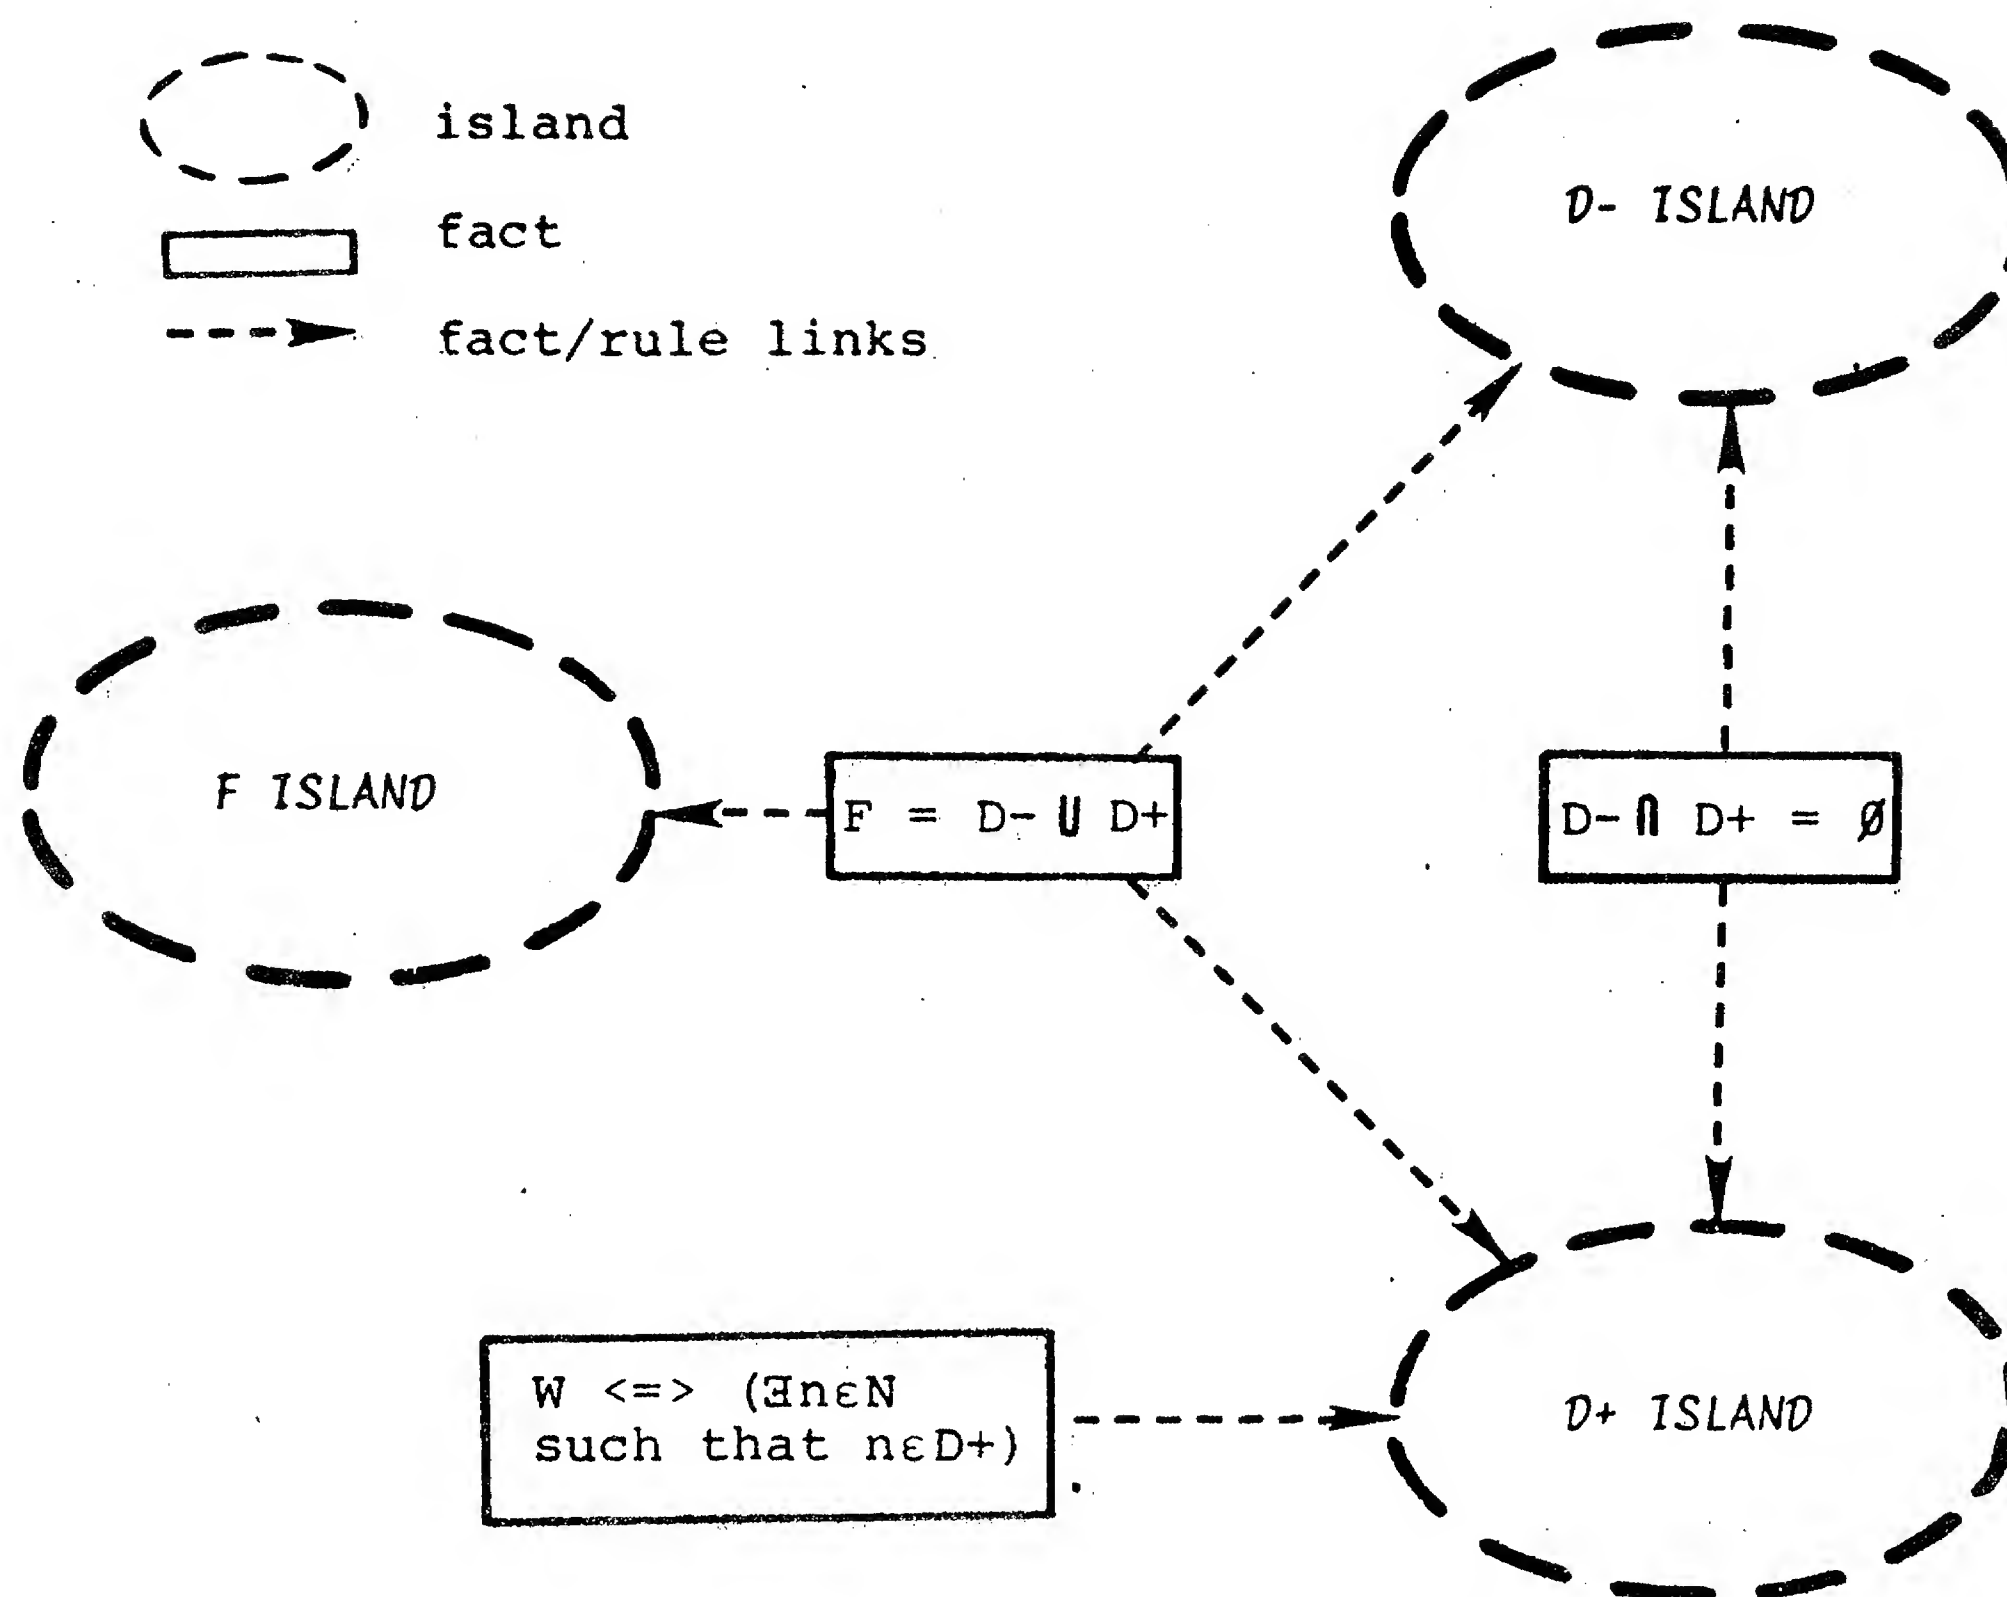
The GG employed in the student simulation testbed has not yet been augmented in this fashion.

This extension will be important because of the possibility that the same evolutionary relationships linking procedural rules can play a role linking declarative statements. One logistic statement can be a generalization of another, or analogous under some mapping of constants, or a refinement. With such an extended GG, the coach could tutor both procedural and declarative knowledge, obtaining leverage by moving between the two in the light of the student's current difficulties.

4.3 The extended GG represents planning knowledge.

Not all knowledge about rules describes their evolutionary relationships. Since sets of rules form problem solving programs for the task, we should expect, as with all programs, that knowledge about their order of application must be represented. There is no difficulty in extending the GG to represent this knowledge. It is only necessary to define the appropriate links. For this

Fig. 6. Islands of Knowledge Have Declarative Foundations

Cave Sets

F = fringe caves
 $D-$ = safe fringe caves
 $D+$ = dangerous fringe caves
 \emptyset = null set
 N = neighbors of the current cave

reason, prerequisite and postrequisite relations are defined.¹⁷ Fig. 5 illustrates various planning relations. For example, a postrequisite link insures that the $D+$ rule "If there is a warning, add the

17. An alternative is to supply meta-rules that specify the order of application. This is a useful approach when the to goal is an expert program, but it is not sufficient for the tutoring context. A meta-rule that specifies that a trio of rules be executed in the order $R1, R2, R3$ does not tell the tutor whether this is the only order or merely one among a set of possible orders. The tutor must know if it is to respond appropriately to the student's idiosyncratic approach. The planning links provide only the basic ordering constraints. Sacerdoti [Sac75] makes a similar argument for planning networks to facilitate self-debugging on the part of a problem solving system.

neighbors to D+." is followed by the rule "Subtract the visited cave set from D+." This second rule is needed since some of the neighbors added to D+ may have been already visited, and hence are safe.¹⁸

Thus, the extended GG incorporates planning knowledge. As with any addition of knowledge, it supplies tutoring leverage. The coach can now expect that for certain rules, the student's difficulty may lie not in knowing the rules, but in understanding their order. The consequences of such confusion can be explored in the student simulation testbed. Below is illustrated a situation where the student has apparently forgotten to execute a prerequisite rule and shows the advice the new GG-based coach might offer in such situations.

The following interaction with the Wumpus game was obtained from a version that allows the player to record his or her hypotheses on the display screen. The player's responses are underlined. The final tutorial intervention (in italics) is a hypothetical one by the GG-based coach now under construction.

We are now in cave D. The neighbors are caves C, G, and B. What now?

> B

We are now in cave B. The neighbors are caves F, D, and A. What a stench! The Wumpus must be in one of the neighboring caves. Squeak! I hear bats. They must in one of the neighboring caves. What now?

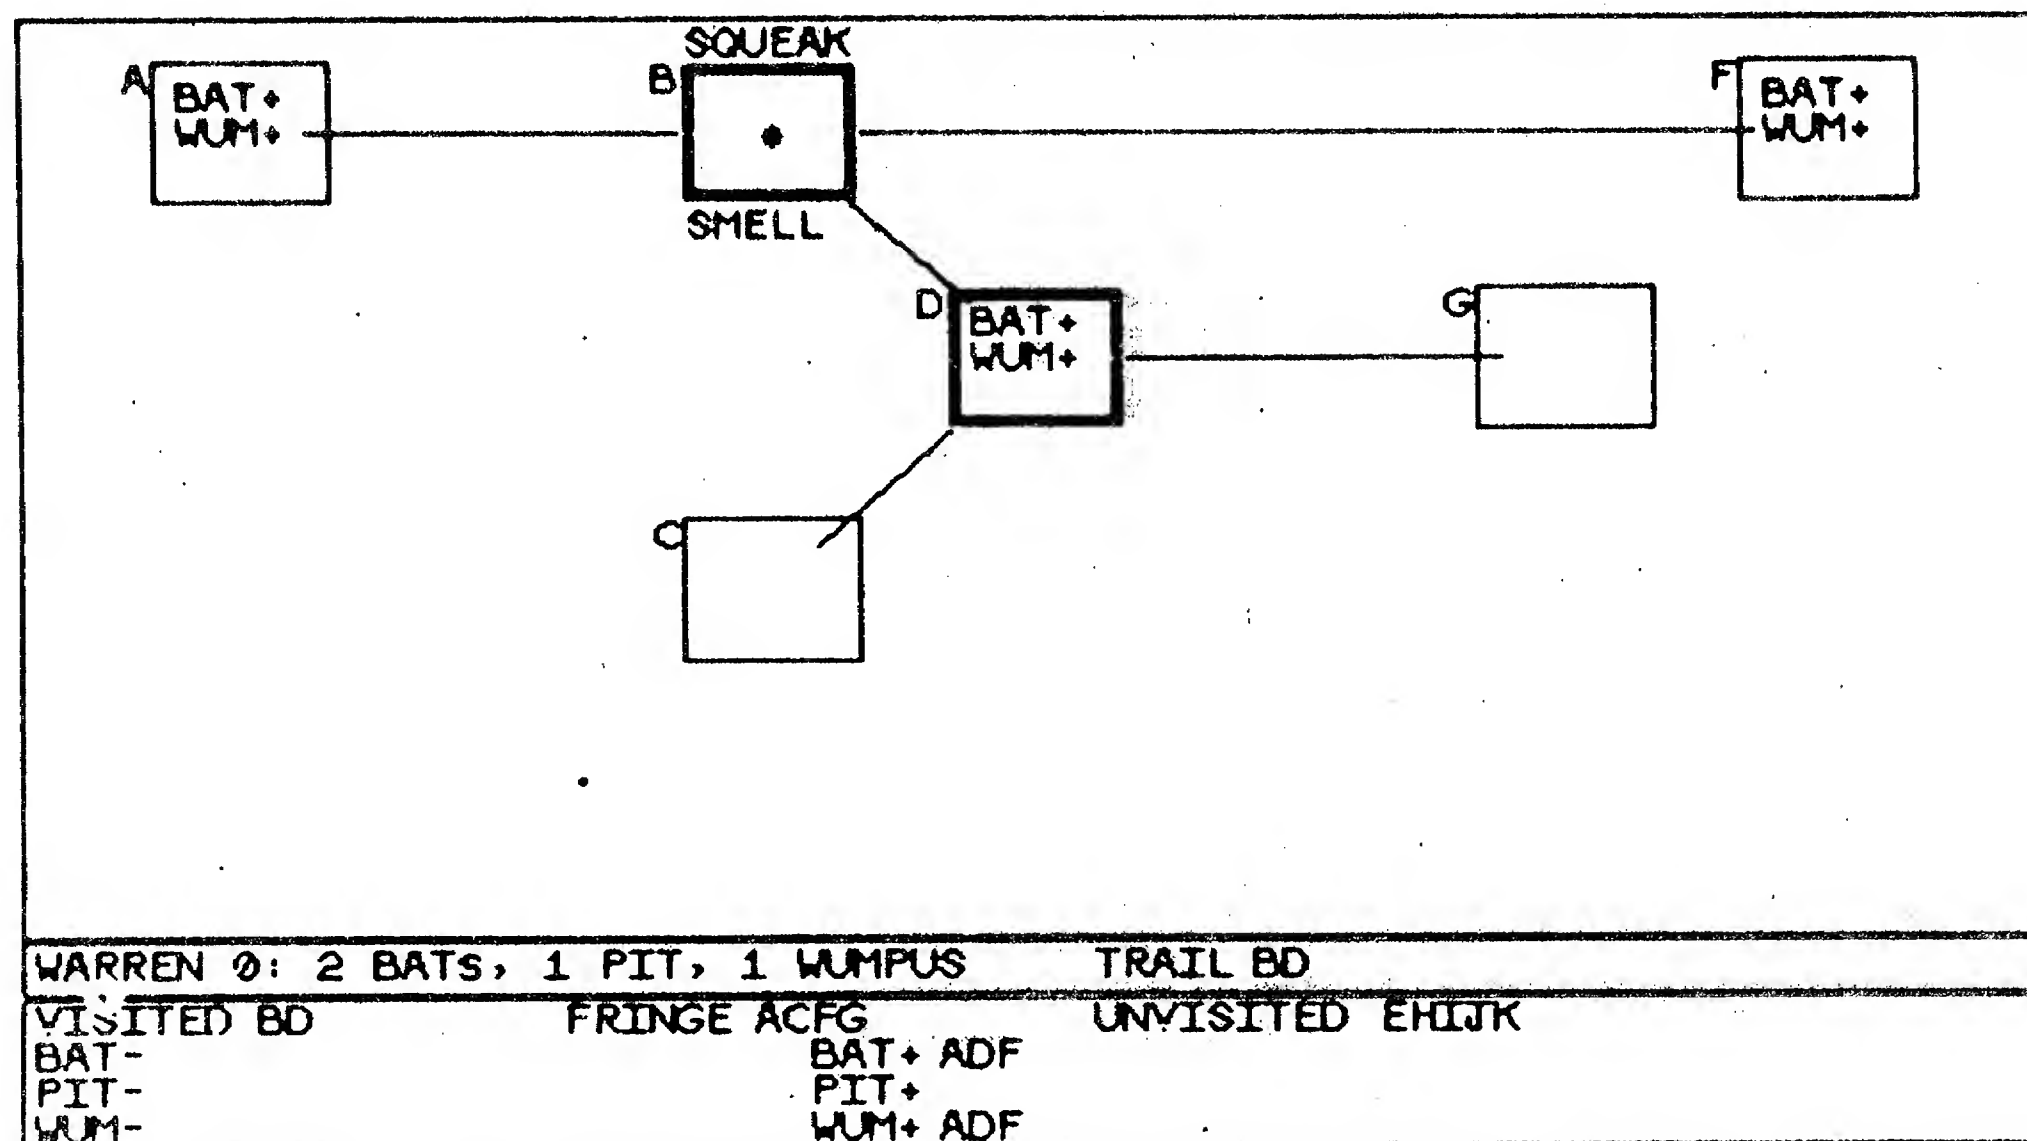
> X+ ; *This command marks caves that may contain a danger.*

Which danger (Bats, Pits or Wumpus)? BW. Which caves? AFD
; *The result is that the display screen shows a warren in which BAT+ and WUM+ markers appear in the designated caves. (Fig. 7 is a snapshot of the display in this state.)*

What now?

18. A single rule could have been written: "If there is a warning, add N-V to D+", but breaking this procedure into two rules allows a finer grain of modelling and tutoring. The coach must be able to identify the deviant simplification of adding the neighbors without pruning the caves already visited. Thus as a general philosophy, rules are broken into small chunks with planning knowledge made explicit via links between rules.

Fig. 7. The Warren Display after the Dangerous Caves are Marked



> C

Ira, you have correctly recorded that caves A, F and D risk bats based on the squeak in cave B. However, you have moved to cave C without concluding that cave D is safe. You know cave D is safe because you have visited there. You might want to record this knowledge to avoid confusion.

4.4 The extended GG is a network of islands.

In conclusion, the extended genetic graph is not different in kind from the basic genetic graph. It remains a graph of knowledge nodes linked by various genetic relationships. Its increased power derives from a hierarchical structure of grouped rules (islands), an extension of individual nodes from representing only rules to representing both rules and facts, and an augmentation of the link set to include control knowledge.

Implicit in this structure is the following view of learning: new rules are constructed from old in terms of processes corresponding to the individual links. However, the graph does not describe a unique evolutionary path. One learner may rapidly acquire a generalization, another may first

build several specializations before constructing the generalization, while a third may never acquire the generalization. Hence, the tutor should encourage this idiosyncratic construction of new knowledge by giving advice appropriate to the learner's current knowledge state (position in the graph) and particular style of learning (preference for particular links). The redesign of the Wumpus coach to employ the guidance of the GG to more closely approximate this ideal tutoring behavior is the subject of the next two sections.

5. The genetic graph is a basis for tutoring.

The GG guides the Tutor component of an AICAI system in two ways. First, it suggests which skills to discuss with the student, namely those at the frontier of the student's position in the graph. Second, once a skill is chosen for discussion, the GG supplies guidance for explaining that skill in more than one way by means of relating it to its evolutionary predecessors.

5.1 The genetic graph suggests the tutoring topic.

In Script-based CAI, the order in which topics are introduced is predefined. The student proceeds to the next author-supplied question after he has successfully answered the current query. This has the advantage that the author can control the introduction of material in the light of his understanding of the subject matter, but the disadvantage that the order is rigid.

Expert-based CAI is less rigid since it has the power to allow the student to explore a problem in his own fashion, analyzing his responses in terms of an underlying skill set. Tutoring is oriented around supplying advice in those situations wherein the student has chosen a less than ideal option. But the Expert-based tutor has no guidance with respect to whether discussion of a given skill is premature in the context of those skills the student has already acquired.

Providing a genetic graph addresses this limitation. If we accept the educational heuristic that learning is facilitated by being able to explain a new skill in terms of those already acquired, the skills with the highest priority for being taught are those on the "frontier" of the student's knowledge model. Employing this heuristic, the AICAI tutor can limit its intervention to those situations with "leverage", namely those that involve the discussion of a skill on the frontier. For example, consider two students: a beginner who has mastered the basic fringe rules and whose frontier is the dangerous and safe islands, and an intermediate player who has mastered these islands and whose frontier is now at the multiple evidence island. Let us consider what kind of tutoring the AICAI coach should offer if the player moves to cave 14 in the scenario game of fig. 2. Recall that 14 is a bad choice because the existence of double evidence makes it likely that a fatal pit is there. For the intermediate student, the tutor would intervene -- there is leverage to describe the double evidence heuristic in terms of its evolutionary predecessors; for the beginner, the tutor would not -- there are no available genetic links with which to build an explanation.

The GG does not solve the "choice of topic" problem. It offers the frontier as a preferred subset, but the Tutor must still choose among this subset or possibly decide to reject it entirely.¹⁹ To make this decision, the Tutor must apply general teaching heuristics ("Vary the topic discussed!") and student specific strategies ("Maximize the opportunity for 'discovery learning'; that

19. There are alternatives to frontier tutoring. A teacher could seek to explain the syllabus as a whole to give the student perspective. Then later the teacher could return to a given subset of the syllabus and refine the student's knowledge. Norman refers to this approach as Web tutoring [Nor76]. It is more useful however for a syllabus of facts than one of procedural skills. The reason is that skills have prerequisite relations that prevent advanced skills from being used before simpler ones are acquired. Static facts don't generally have such a rigid ordering. Thus our skill tutors usually do not employ the Web technique.

Nevertheless, it is possible to explain a skill whose evolutionary predecessors have not been acquired by constructing a long explanation *ab initio*. The frontier heuristic biases the system against such an approach, but the Tutor may be required to employ it in some situations (the frontier skills have already been explained many times and the student appears to need some perspective on the syllabus) and for some kinds of syllabi (the skills are largely independent of one another).

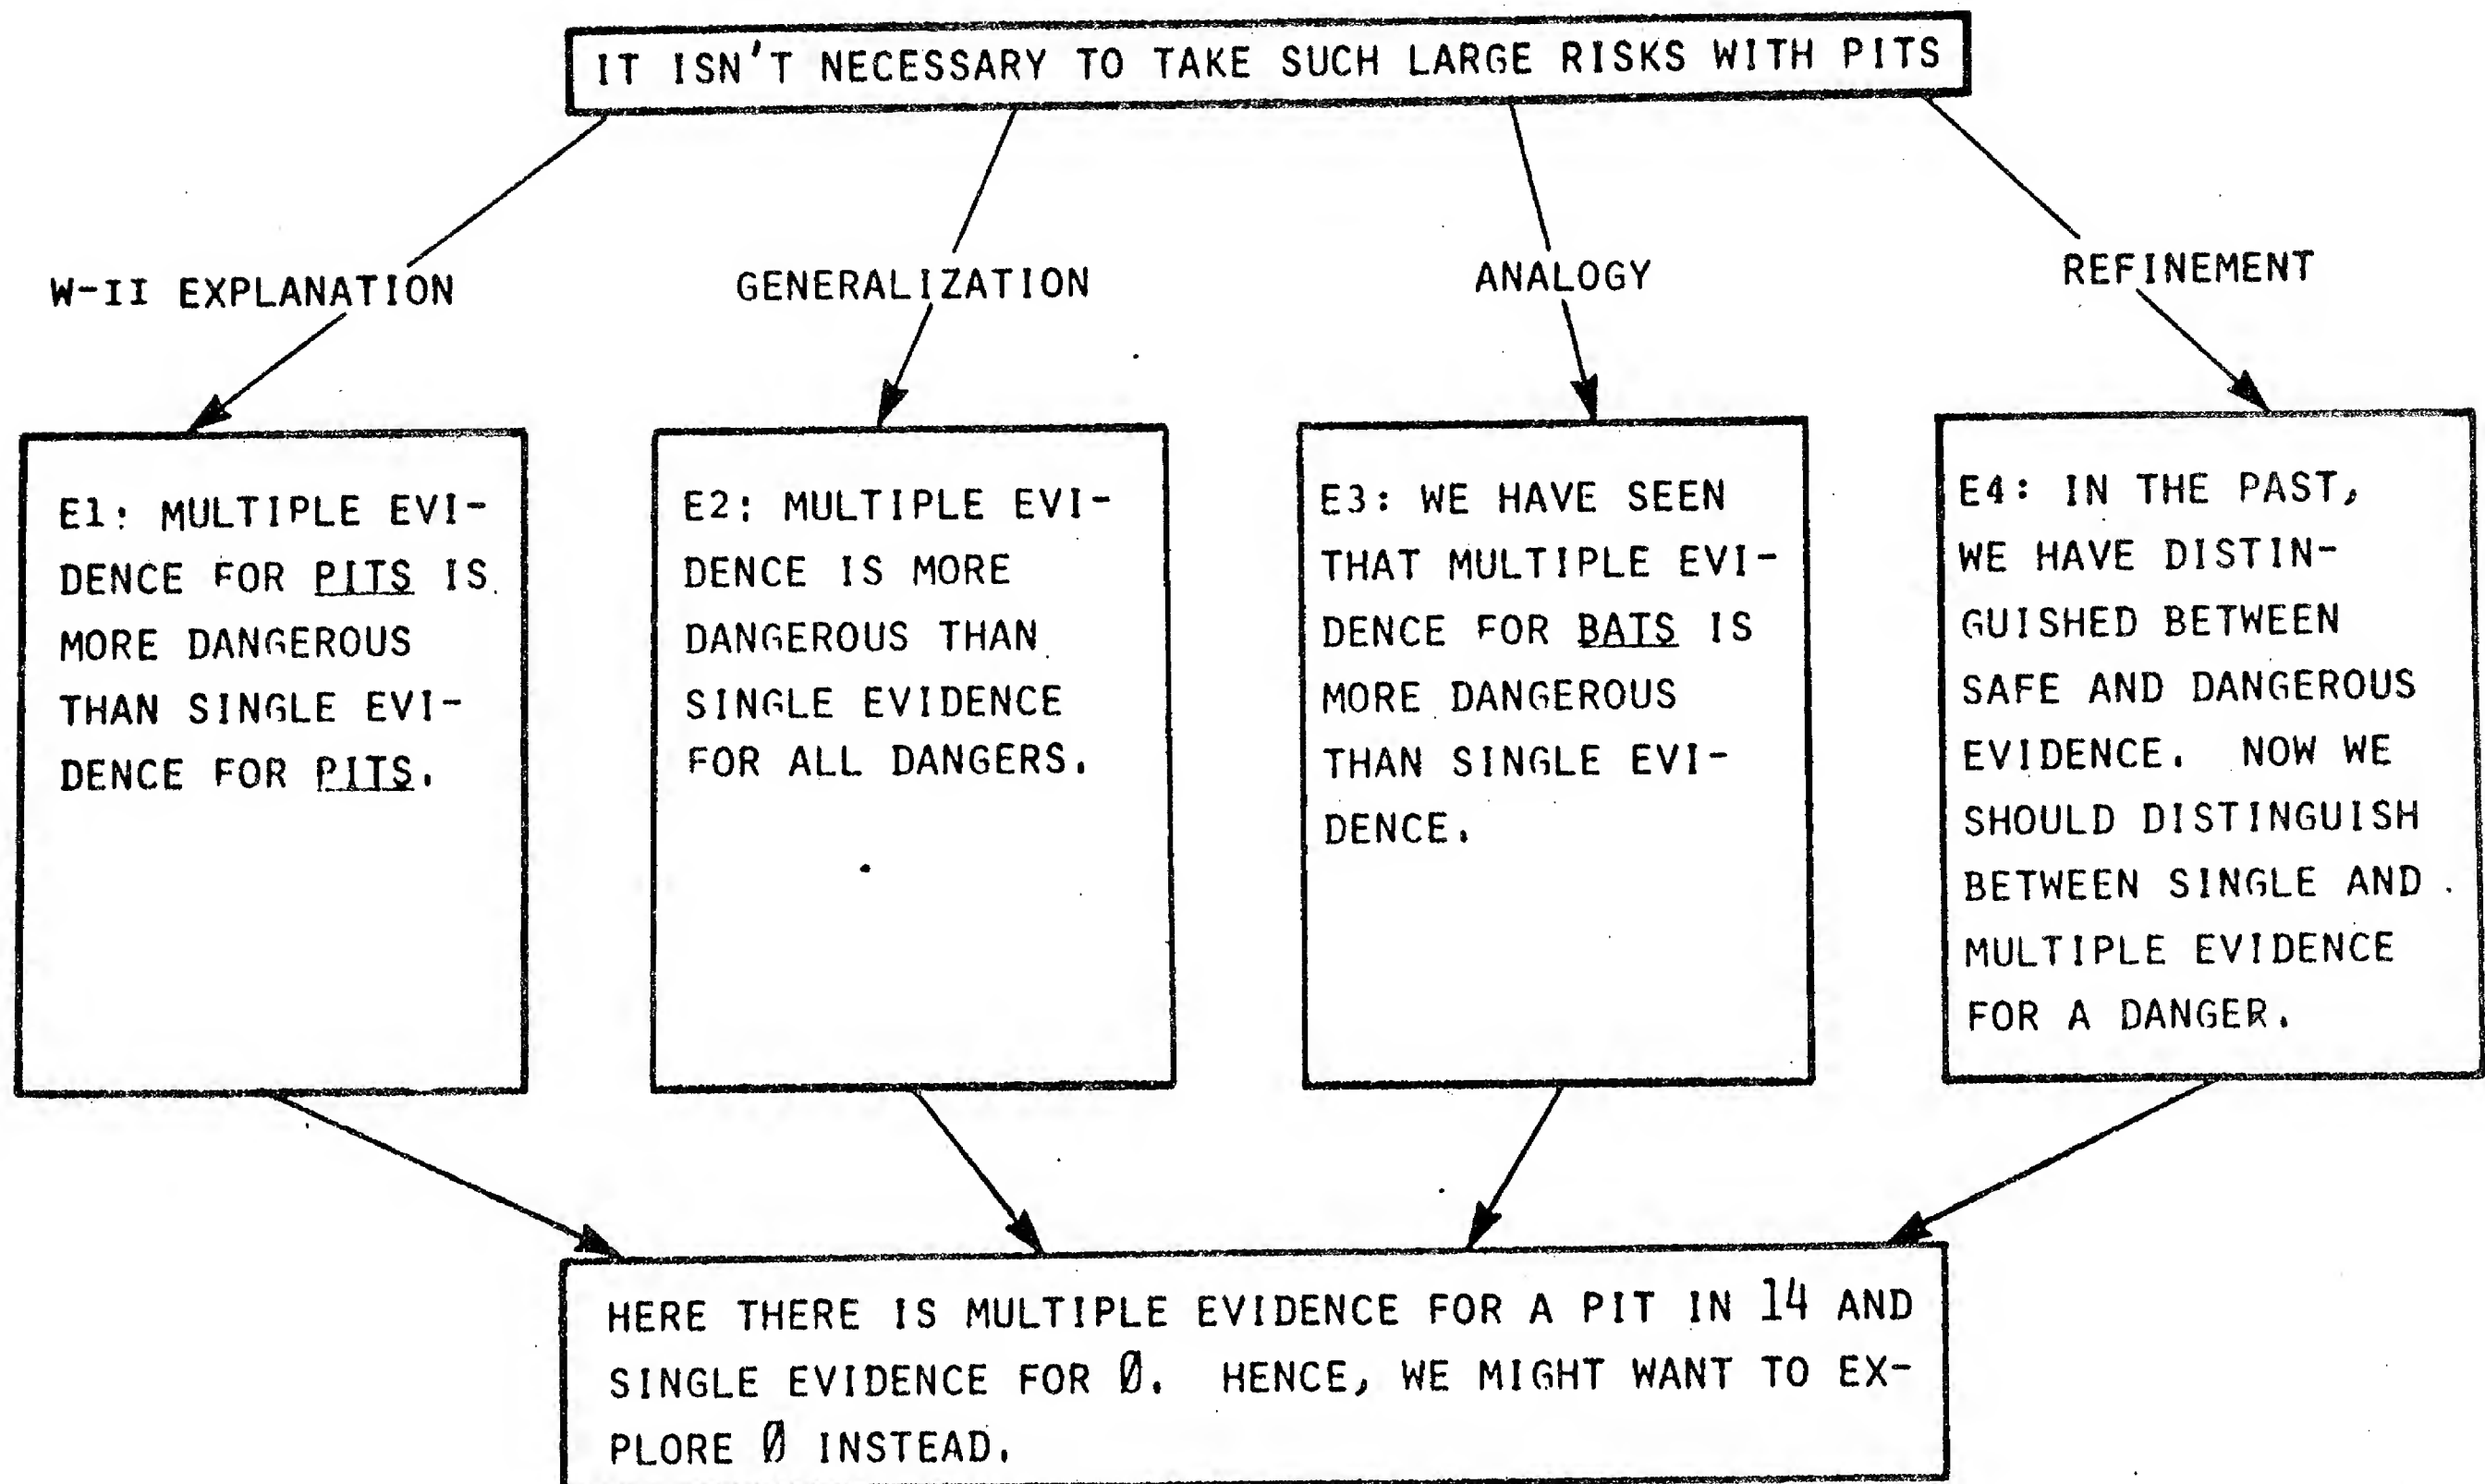
is, do not discuss any topic at all when the student's progress through the syllabus is proceeding at a satisfactory rate!"). The role of the GG in this context is simply to make available to these teaching heuristics the epistemological relations between the skills of the syllabus.

5.2 The genetic graph supplies multiple explanations.

Once a topic is selected, the ability to explain that topic in more than one way is an important tutoring technique. Script-based CAI achieved this explanatory power by supplying "author" languages in which clever explanations could be written by teachers. Expert-based CAI, by eliminating scripts, lost this power. But in return it acquired the ability to respond to a larger number of situations, albeit by means of a restricted number of machine-generated explanation types. Genetic AICAI retains the Expert-based CAI ability to respond to a large number of situations, but adds the capability to explain a particular skill in diverse ways. This capability derives from the ability to explain a new rule in terms of its genetic links. For each link type, the tutor is provided with an explanation strategy. For example, fig. 8 shows three variations on a Wusor-II explanation generated by explaining the "avoid multiply dangerous situations" rule in terms of its evolutionary relatives. (The basic WUSOR-II explanation is the one we examined earlier for the poor move to cave 14 in the game state of fig. 2.)

As with the selection of the rule to be discussed, the choice of explanation for that rule is not determined by the GG. That choice depends on general teaching heuristics (such as "Vary your explanation!") and student specific criteria (such as "Avoid strategies which have been consistently unsuccessful in the past!"). The role of the GG, however, is to increase the available choices on which these selectional heuristics operate.

Fig. 8. Variations on an Explanation



5.3 Tutoring using an Extended Syllabus Representation

The range of tutoring strategies is increased further by employing the extended syllabus described in the previous section. Examine again fig. 5. The islands of rules for dangerous and safe caves are linked by a "bridge analogy," a generalization of the analogy linking the individual rules. By tutoring the bridge analogy, explanatory leverage is gained by providing support for an entire group of rules in a single explanation.

Similarly, a declarative foundation provides another opportunity for generating support for an entire set of rules. Here the common link is from a set of declarative facts to the island of rules they imply. For example, the tutor, using the declarative foundation, might discuss the general importance of the concept of multiple evidence rather than the specific rules employed to maintain

the single and multiple evidence sets. Again, the tutoring hypothesis is that by discussing the declarative foundation, the student will deduce a group of related rules on his own. It is therefore a potentially powerful tutoring strategy.

5.4 The genetic graph does not solve the tutoring problem.

Tutoring is a complex task for AI-based CAI systems that do not have access to author-supplied scripts. The system must decide (1) whether to intervene, (2) what topic to discuss, and (3) how much to say about that topic.²⁰ The GG does not decide these questions. However, it does serve, first, to constrain the set of topics by defining a frontier, and, second, to extend the variety of explanations available for discussing the topic of choice.

6. The genetic graph is a basis for modelling.

To offer appropriate tutorial advice, a teacher must accurately model the student. The GG facilitates the modelling process in an AICAI tutor in three ways. First, the nodes of the graph provide a more refined structure for a model of the student's knowledge state than the skill sets of subset AICAI systems. Second, the organization of the graph provides a metric regarding which skills the student can be expected to acquire next. Third, the links of the graph provide a complementary structure for a model of the student's learning behavior.

20. See [Col75] for a study of Socratic intervention strategies.

6.1 The student knowledge model overlays the nodes of the genetic graph.

Script-based CAI systems build student knowledge models by maintaining statistics on the correctness of the student's answers. The validity of such models is severely limited by the restricted capability of the script to judge correctness, having only a list of expected responses on which to judge the answer.

Expert-based CAI systems escape the limitation of the script by constructing their student knowledge model from hypotheses regarding which skills of the *embedded expert* the student is believed to possess. I have termed such models "overlays" [Gol77a] to emphasize that their structure is derived from the structure of the underlying expert system.

As an example of the improvement of expert-based modelling over scripts, consider Wumpus. The embedded expert of the WUSOR-II coach can evaluate any game state that arises. The number of such states, given an arbitrary number of caves, of dangers in these caves, and of student paths through the resulting maze, is enormous. Scripts of correct answers are clearly out of the question.

But expert-based models have a fundamental limitation. They fail to consider that the novice student may not be employing a subset of the expert's skills, but rather using simplifications, deviations, and other evolutionary predecessors of those skills.²¹ Given our GG, the extension is clear. The student's knowledge model will be constructed as an overlay, not on the final set of skills, but on the GG itself.

21. In certain situations, there is a rationale for expert-based models. The "expert" may be one selected to be only minimally in front of the student. Or the task may be sufficiently restricted that novices are generally subsets of the expert's skills. Or the skills themselves may be broken into small "micro-skills" so that modelling in terms of the presence and absence of these micro skills is reasonable. Indeed, the Genetic AICAI system reduces to the expert case if the GG does not in fact contain other than a single subset of skills. Thus, the expert-based CAI system can be profitably viewed as a simplification of the Genetic AICAI system.

6.2 The genetic graph guides the construction of the model.

Given the form of the model as attributing regions of the GG to the student, it is now appropriate to examine how the model is induced. I shall describe the basic method employed by expert-based CAI programs, and then construct an improvement based on the learning metric implied by distance between skills in the GG.

Expert-based CAI constructs the student knowledge model by hypothesizing that a student does not possess a skill if the student's answer for a given situation is worse than the answer the expert could deduce based on that skill.^{22,23} To illustrate this, consider again the scenario of fig. 2. If the student chooses cave 14, which is more dangerous than its fellows by the multiple evidence skill, WUSOR-II increases the weight of its hypothesis that the student does not possess this skill.²⁴

This method of comparing embedded expertise to student performance remains basic to the Genetic AICAI system, but is improved as follows: the GG is viewed as defining a number of "players" of increasing power, corresponding to intermediate skill plateaus in the graph. For Wumpus, there are five such players defined in terms of the five phases of Wumpus skill.

22. And contrariwise, if the student chooses the expert's choice, then the coach hypothesizes that the student is familiar with those skills the expert employed to determine that the move chosen was best.

23. In fact, the process of modelling is more subtle than this. For each situation analyzed, the raw data is recorded as increments to two variables associated with each skill: APPROPRIATE which records how many times the Expert believed the skill was appropriate and USED which records how many times the player was believed to have employed the rule in appropriate situations. Their ratio forms the FREQUENCY of use of the skill. The AICAI tutor acts as though the student knows the rule when this ratio exceeds a threshold. The complexities in maintaining such a model are discussed in [Carr77b].

24. This simple modelling method is improved by the capability in some AICAI programs to take account of the student's background, and in some situations, to ask the student explicitly why he chose a certain option. These improvements are orthogonal to the improvement the GG allows in the fundamental method. They are discussed in [Carr77b].

Phase 1: rules for visited, unvisited and fringe caves.

Phase 2: rules for possibly dangerous, definitely dangerous and safe caves.

Phase 3: rules for single versus multiply dangerous caves.

Phase 4: rules for "possibility sets", i.e. keeping track of the sources of dangers.

Phase 5: rules for numerical evidence.

Each of these "players" examines the student's move and proposes which skills the student appears to be employing. These hypotheses are attached to nodes of the GG. The overall belief that the student possesses a given skill is a summation over the hypotheses of the individual players.

If it were the case that the student might possess skills from anywhere in the GG with equal probability, then all of these players would have equal weight when formulating the overall hypothesis. But, the GG embodies a theory of the evolution of the learner's knowledge. This theory is just that knowledge evolves along genetic links -- from simplification to elaboration, deviation to correction, abstraction to refinement, specialization to generalization. For that reason, the hypotheses generated by advanced players further and further away from the current plateau are assigned less and less weight.

The result is a desirable conservatism in the modelling process. This is reasonable, since it accords with the common sense educational heuristic that a radical improvement in the play of a student is more likely due to luck than a discontinuous jump in his skills. By the same token, a radical deficiency in a particular move is more probably due to carelessness than a discontinuous jump to some earlier knowledge state.

This conservatism does not prevent the AICAI coach from ever believing in discontinuous jumps in the student's knowledge. Those players based on skills far from the student's current position in the graph are given some weight. Hence, the coach will eventually accept a radical change in the student's knowledge. But the conservatism is important: without it, the coach has no capability at all to observe the lucky guess or occasional careless move. Hence, the metric on learning defined by the GG supplies a stability missing in expert-based CAI systems.

6.3 The student learning model overlays the links of the genetic graph.

There is still a third dividend to the GG: its links provide the structure for a learning model. In the previous section, we discussed the coach's ability to explain a rule in multiple ways based on the various genetic links associated with that rule. Now given a Knowledge model, the coach is in a position to observe the effect of a given explanation type. It can determine whether the student employs the skill in subsequent play. If a given explanation strategy consistently leads to skill acquisition, it is reasonable to believe that this explanation strategy is a successful one for the particular student. If not, then the opposite hypothesis can be induced, i.e. that the explanation strategy is not a successful teaching strategy for the particular student. Thus, a *Learning overlay* can be generated over the set of genetic links that maintains a record of the effectiveness of the explanation strategy associated with that link type.

The use of such a model is straightforward: it serves to personalize the choice of explanation strategy for a particular student by selecting from those that have proven successful in the past.

6.4 The genetic graph does not solve the modelling problem.

Constructing a model of the student's knowledge and learning attributes is a complex task for a human teacher to perform. It is certainly the most difficult activity of an AICAI tutor. The genetic graph provides a framework for this modelling. The student's knowledge is described in terms of the nodes of the graph; his learning behavior in terms of the links; his progress in terms of paths in the graph. It provides a more powerful foundation for modelling than either a script of correct answers or a set of expert skills.

Nevertheless, the GG does not solve the modelling problem. While the process of constructing a model gains guidance from the graph, it remains complex. No particular answer by the student is certain evidence. He may have misunderstood the question, or lost interest in

formulating an answer, or changed his goals entirely. The coach, given its inability to observe the student's facial expressions, understand his language, or indeed even know whether he is at the console thinking or simply taking a stroll, is at a severe disadvantage compared to a human teacher. And modelling the student is among the most difficult tasks for skilled human teachers. I term this the "bandwidth problem". No matter how excellent the GG is as a representation of the knowledge being acquired, modelling is dependent on observing this acquisition. Hence, methods for increasing the bandwidth with which the computer coach can observe the student are an important supplement to the GG in model building.²⁵ The virtue of the GG is simply to provide a target data structure for the evidence gathered by this increased bandwidth.

There is another deeper limitation to the modelling paradigm offered here. While it is true that one can only model what one understands, it is not true that one must represent the syllabus in such an explicit form. A human teacher can be expected to grow his understanding of the task in response to observing the student's behavior. For the more general situation of tutoring in large open-ended worlds, this is necessary; however, it involves the incorporation of a learning capacity into the coach, a non-trivial though important function. The next section discusses a preliminary formulation of the learning theory that would be required.

25. For Wumpus, we are currently exploring several kinds of "assistant programs" that serve to increase the bandwidth with which the Coach can observe the student. One assistant offers the display screen as an interactive medium to replace the pencil and paper the student uses to draw the warren and record his hypotheses. In this fashion, the coach can observe that part of the student's intermediate reasoning that is overt. It is our expectation that this graphic assistance will make a major improvement in the accuracy of the Coach's model.

7. The genetic graph is a basis for learning.

Implicit in the genetic graph is a theory of learning. This section explores this theory and considers its implications for the design of computer coaches. The model of the student suggested by the genetic graph is shown in fig. 9. The processes of the student are divided into two homunculi²⁶ -- a problem solving specialist and a learning specialist -- with the graph serving as the student's basic memory structure for procedural knowledge. The problem solving homunculus applies the program defined by the frontier of his genetic graph to the current task. The learning homunculus extends the genetic graph in response to new tasks, tutorial advice and observed difficulties of the current program.

The learning homunculus consists of a set of strategies corresponding to the various links of the graph. Its task is to build new rules, leaving behind -- as a record of its operation -- links which connect the new rules to their evolutionary predecessors. The links are labelled with the learning strategy responsible for the construction.^{27,28}


The genetic graph offers only a structure for a learning theory. It suggests that the learning processes consist of procedures which generate the various links, but it does not describe the details of these processes. It does not enumerate what criteria are used to form analogies, recognize

26. I use the term "homunculus" to emphasize that the learning and problem solving components are envisioned to be machines of exactly the same power. Their only difference lies in their programs.

27. The links are left behind because they themselves can serve as input to the learning strategies. The existence of a profitable analogy can suggest that more analogies "of an analogous kind" are possible. For example, an analogy between the rules of bats and pits can suggest a similar analogy between bats and wumpii. It may not be exact, but the suggestion offers a direction for the learning homunculus to explore.

28. It is of course a simplification to believe that the entire genetic graph remains available to the learner. In fact, there must be a process of forgetting. This process must exist partly to avoid an indefinitely growing use of space and partly to eliminate outdated knowledge that would serve only to misguide the learning processes. A theory of forgetting is crucial to an overall theory of learning and of teaching, but goes beyond the scope of this paper.

Fig. 9. Homunculus Model of the Student



deviations, induce generalizations, or construct conceptual refinements.²⁹

However, this structure is of use, for it focusses our attention on issues involving the interaction of the teaching and learning processes. Four of these issues which I discuss below are: (1) the student as an active agent, (2) a genetic graph for learning, (3) a theory of belief, and (4) the topology of the graph as a measure of learning complexity.

29. Enumerating such criteria has been the focus of much work in AI, including Winston [Win75], Evans [Eva68], Moore & Newell [Moor73] and Richard Brown [Brow77] on analogy; and Sussman [Sus75], Goldstein [Gol75], and Sacerdoti [Sac75] on debugging.

7.1 The student is an active agent.

The model of the student presented above emphasizes the viewpoint that the student is an active agent, engaged in a constructive process of generating new knowledge. From this perspective, the tutor's objective is to encourage this process in the student. This reminds us that the current activity of most AICAI tutors -- intervening and supplying a complete explanation -- is only one end of spectrum of tutoring activity. At the other end of the spectrum is "tutoring without talking", that is saying nothing at all, but instead altering the problem domain in order to facilitate the learning process.³⁰

There are as well a range of intermediate interventions between these two extremes. An example is that the tutor could suggest that a rule exists that could be applied in the current situation which is analogous to some already acquired rules, but not specifying the new rule or stating the analogy. The next generation of AICAI tutors should be able to supply advice across this spectrum, altering the nature and extent of their intervention in relation to the current state of the student model.

7.2 A genetic graph for learning skills is possible.

Dividing the student into a Learning Homunculus and a Problem Solving Homunculus raises the question of whether the skills of the Learning Homunculus can themselves be represented as a genetic graph. If there are a collection of rules that define the processes of analogy, generalization, debugging and refinement which are themselves related by genetic links, then explicating this graph becomes an important AI/Psychology goal.

30. "Tutoring without talking" is exemplified by one option WUSOR-II can exercise. It can alter the complexity of the Wumpus game by varying the number of dangers, the propagation distances of their warnings, the number of arrows, and the geometric complexity of the warren. WUSOR-II does this in accordance with its estimate of the student's current level of skill.

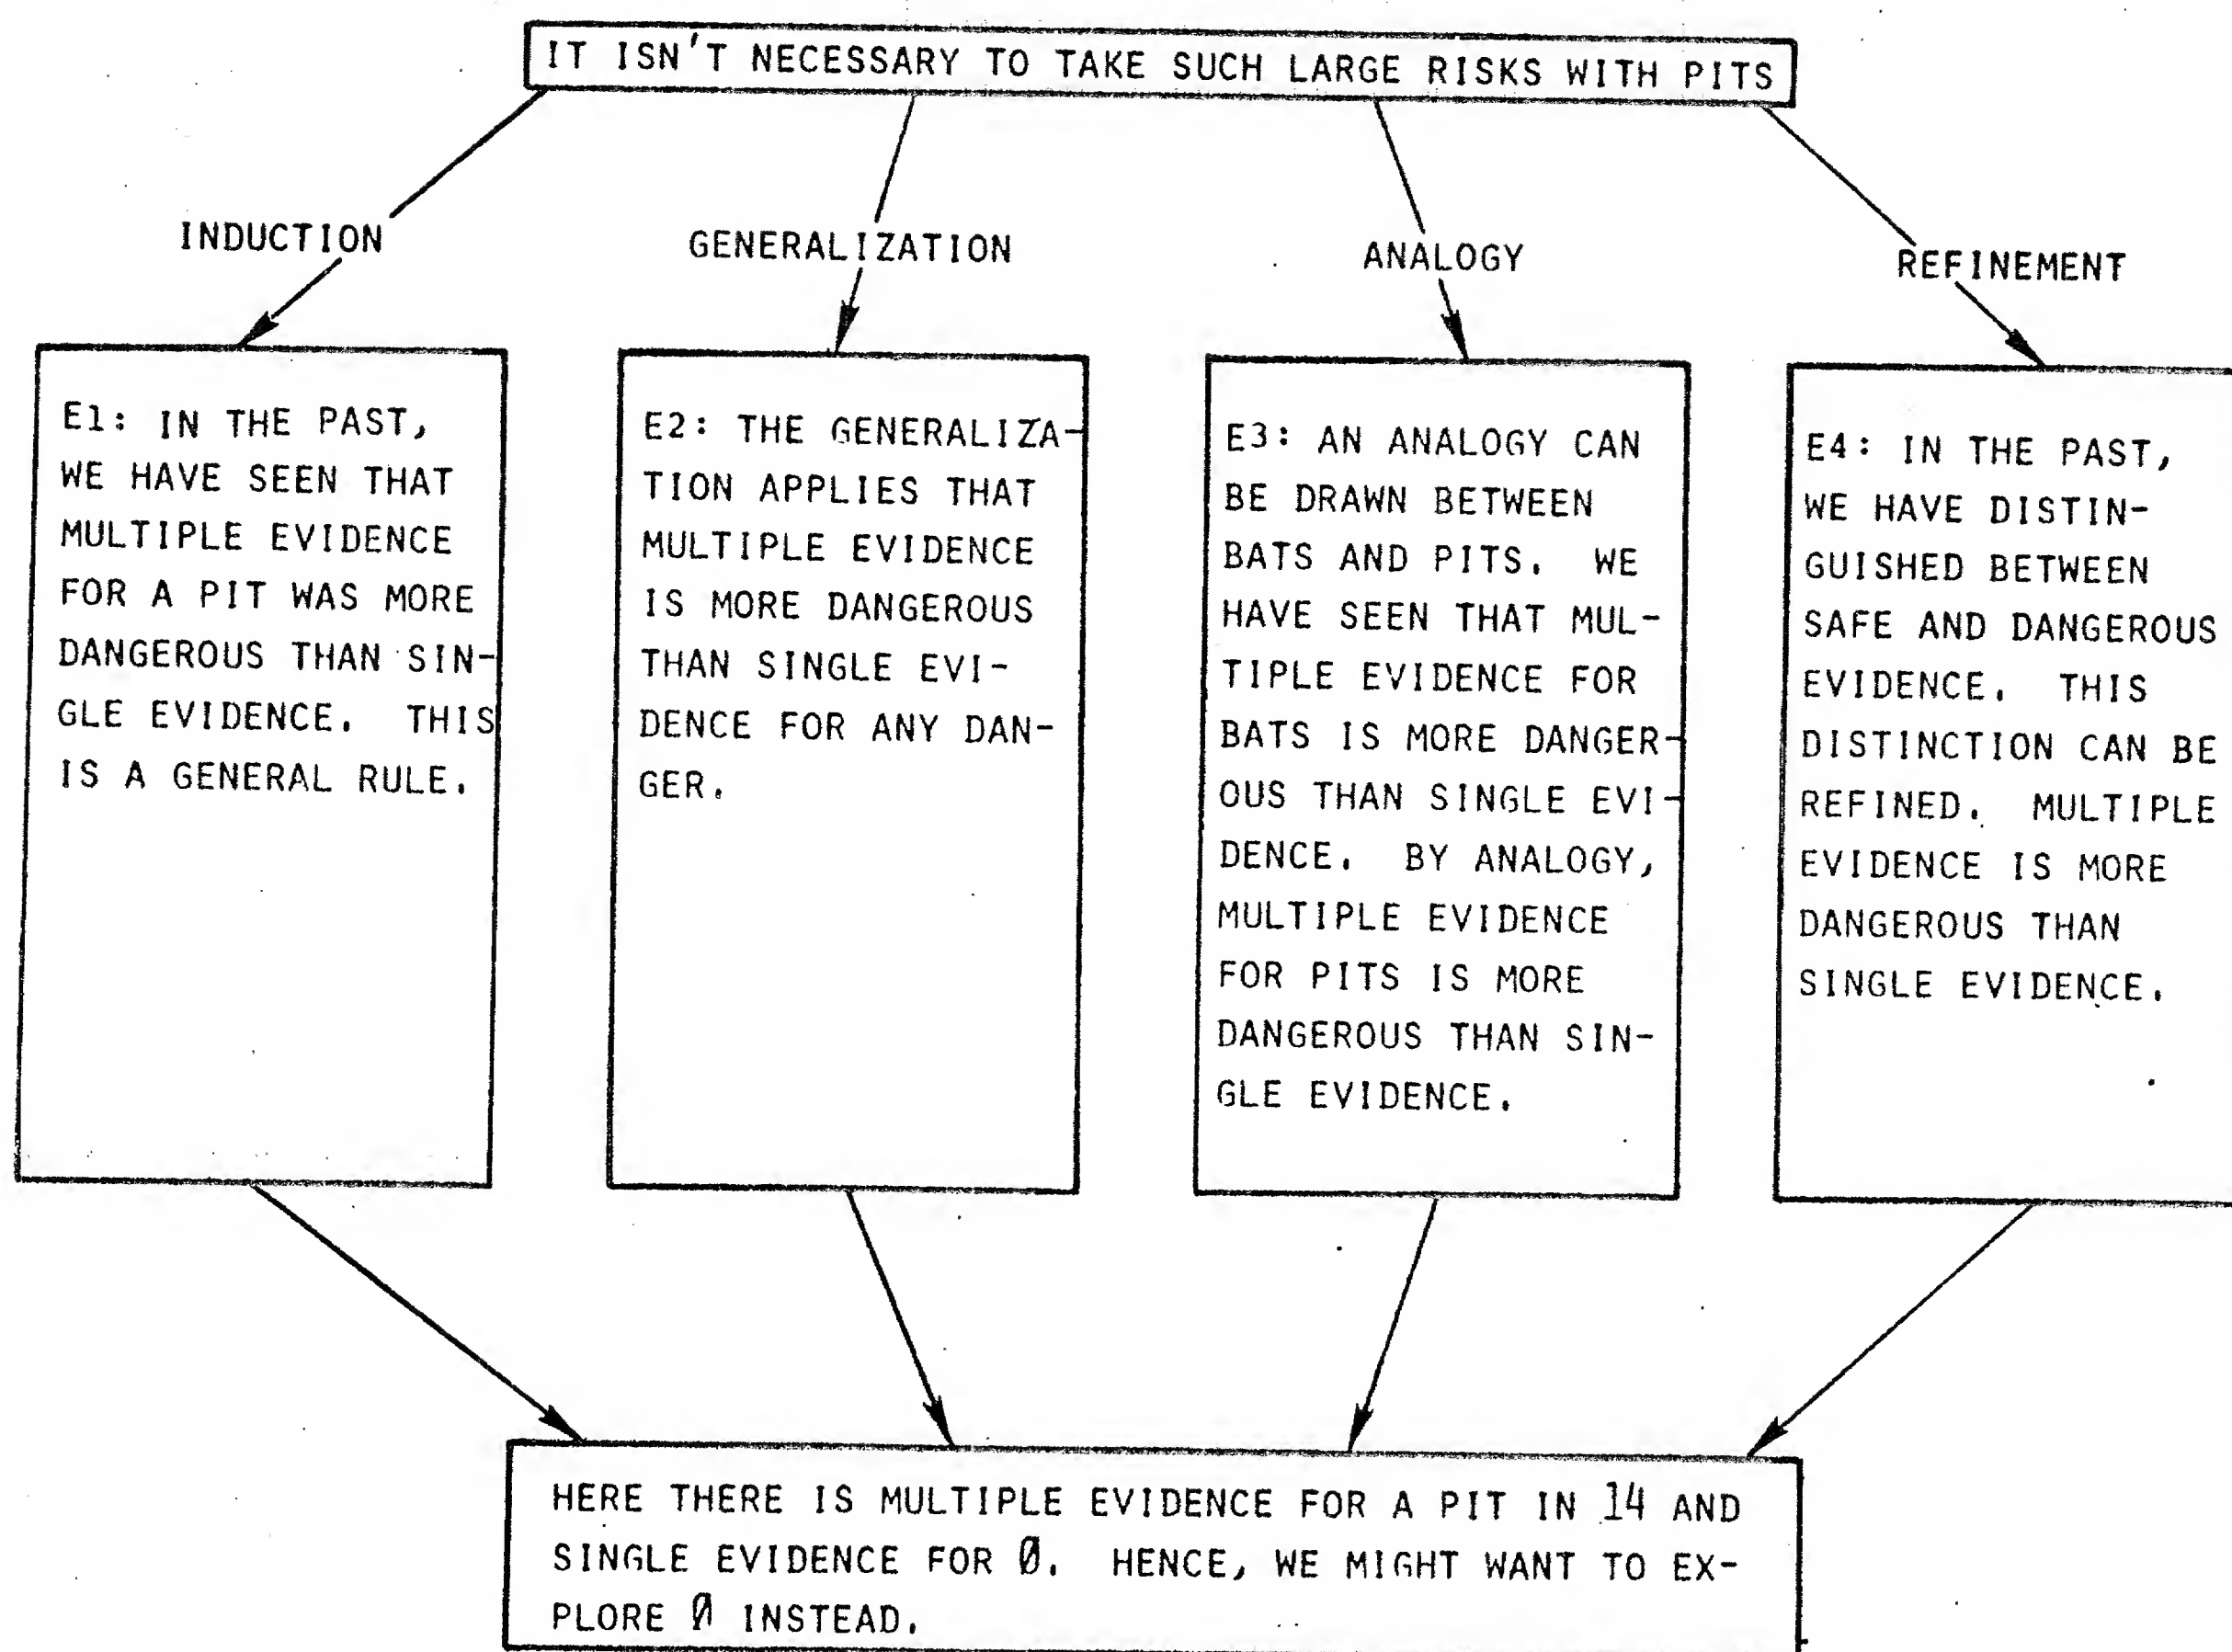
A competing hypothesis is that the learning processes are not related one to another, nor do they have simplifications from which they evolve. They are only an unstructured collection of heuristics, acquired in an isolated fashion. I believe this unlikely, but it may not be simple to explicate a genetic graph for learning.

Constructing a genetic graph for learning skills whose links are again the analogy, generalization and other genetic relationships discussed earlier suggests that still a third L^2 homunculus is not necessary to oversee the acquisition of learning skills. Rather, since the links are the same ones as occur in the domain graph, the Learning homunculus is potentially able to operate on its own genetic graph. Thus the recursion of homunculi is terminated. If this is so, it would be an important result both for artificial intelligence and for psychology, namely that a single learning theory is sufficient for both domain knowledge acquisition and a recursive improvement in the system's own learning capacities.

Naturally for the process to begin, there must be some learning strategies that are innate. Establishing from an AI standpoint which strategies are sufficient to generate the remainder then becomes an important research question.

Given a detailed account of the learning processes themselves, the possibility arises that the Coach might be able to tutor these very skills. As fig. 10 illustrates, its tutoring could be oriented towards pointing out the relevant genetic strategies for constructing new rules. This will be an important direction for future research, since tutoring the skills of any particular domain is less important than tutoring the processes by which these skills are acquired.

Fig. 10. Learning Oriented Explanations



7.3 A belief measure can be defined on the genetic graph.

Presenting both a Learning Homunculus and a Problem Solving Homunculus focusses our attention on the relation between the two: in particular, it raises the question of when a new rule added to the genetic graph becomes a part of the problem solver's program. It is a simplification to speak of the program of the Problem Solving Homunculus being the frontier of the genetic graph. A new rule may represent a misunderstanding, may not be an improvement, or may be as yet incomplete. Hence, some inertia is desirable in a dynamic learning system, if it is not to oscillate wildly or degrade its performance by accepting premature modifications.

This corresponds, perhaps, to the psychological observation that a student does not always employ a skill which has just been explained. While the student may be able to repeat the explanation, and even describe implications of the new knowledge, he may not actually use the skill when solving problems. Teachers recognize this property of students and employ the heuristic of supplying further examples and different kinds of explanations.³¹

A formal representation for this learning conservatism can be added to our learning model by introducing a belief measure. We can restrict a skill on the frontier from being employed by the problem solving homunculus until "belief" in this new piece of knowledge exceeds some threshold, where "belief" is a function of the number, kinds and recency of explanations and examples that have been provided. In terms of our genetic graph representation, we can say that a new rule is not employed until its linkage into the genetic graph is sufficiently strong, i.e. belief in the rule, defined in terms of the number and kinds of links that attach the rule to the existing graph, exceeds some threshold.³²

Such a metric can improve the tutor's expectations about the student's use of a rule following its introduction. The Psychologist module maintains a record of its estimate of the student's belief in a rule in terms of the types of explanations provided, their recency, and their number. When belief is below some threshold, the tutor can expect that more explanations will be needed and that the student will be able to describe the rule when queried, but probably not employ it.³³

31. Authors of script-based CAI systems can incorporate this educational heuristic by supplying multiple exercises and explanations. But the scripts do not provide a theory of where such additional advice will be needed.

32. This is a first order theory. The linkage strength depends as well on the number of situations in which the rule has been explained, the time since these links were constructed, etc. However, this first order theory is sufficient to define some interesting learning complexity criteria which I discuss in the next section, and imply some procedural consequences for the Tutor.

33. This threshold can be dynamically adjusted on the basis of the student's performance.

Given this refined model, we can undertake a fine grained analysis of belief criteria in learners. For example, for some students, many examples of a few links may engender stronger belief than single examples of many links. By examination of the student's performance with respect to the occurrence of such bonds, we can explore the tradeoffs between diversity, repetition and recency. There is also the corresponding AI question of which belief metrics result in a reasonable learning rate, which lead to instability, and which are too conservative.

7.4 The genetic graph topology provides a learning complexity measure.

Focussing on the genetic graph as a record of the learning process suggests a relationship between various topologies of the graph and learning complexity. The utility of this characterization is that it provides guidance to the Tutor regarding which areas of the syllabus require more attention and to the Psychologist with respect to which skills the student can be expected to have difficulty with.³⁴

From a learning viewpoint, the complete genetic graph of the tutor is a roadmap. It describes various paths the student's learning process might take. If the tutor's graph shows that a given rule has many links, then the expectation is that the student will have little difficulty in acquiring that rule himself. There are many opportunities for him to do so. But if another rule has but one link to the other rules, or indeed none, then here is a topology that suggests the need for tutoring advice.

For example, fig. 3 showed a cluster of rules densely connected by generalization and analogy links. Our belief metric suggests that such clusters are easier to acquire than sparsely connected

34. Traditional epistemology discusses validity, not complexity. This is because complexity is not well-defined except in relation to a particular learning theory. Traditional epistemology did not have such a theory. We are developing a theory of knowledge that is not independent of the "knower".

regions of the graph. The procedural import of graph density is to cause the Tutor to expect that repetition will be little needed in dense regions but strongly demanded in sparse areas.³⁵

Thus, topologies of the syllabus suggest a theory of learning complexity. Experiments are needed to determine if this is borne out. But if so, it is an important theoretical idea for education, independent of the use of computers.³⁶

7.5 Designing simulated students is a research methodology.

We intend to explore the many issues raised here by extending our "student simulation testbed" to include computer students which learn. Such students can be used to explore the effect of different belief metrics on stability and of different learning strategies on the growth of the graph.

35. Recall that in our discussion of the genetic graph and its relation to the Psychologist, I introduced a learning complexity metric. This metric was employed to make the Psychologist conservative in its belief that the student's behavior had exhibited a particular skill when that skill was far from the frontier of the student's current knowledge state. Formally this took the form that the K model "appropriate" and "used" parameters are altered proportional to how far the skill is from the student's current knowledge frontier.

The learning complexity implied by the belief metric for certain syllabus topologies suggests a refinement of this complexity metric, namely that sparsely connected nodes should be expected to be more difficult to acquire than densely connected ones, if at the same distance from the knowledge frontier. In particular, with respect to skills on the frontier, the Psychologist should be conservative in believing that a student has acquired a particular skill when that skill is weakly linked to the student's knowledge frontier.

36. It is conceivable that formal analysis of a syllabus with a genetic graph may serve a useful educational function by predicting the learning complexity of the material. If the graph is largely a chain of rules, we can expect difficulty in convincing the student to employ these skills. Their support will rest entirely on repetition of a single explanatory method. On the other hand, if the GG contains many islands, bridges, and clusters, then we can expect that little tutoring may be required due to the rich interconnectedness of knowledge in this domain.

The validity of the formal analysis is not yet established. But its importance is clear. Education rests on at best a pop epistemology. Philosophic epistemology is too removed from learning. If our analysis provides a middle ground, rigorous, objective, and concise but still about the learner's relation to knowledge and not some abstract definition of truth, then we have made progress in developing a theory of education.

Ultimately, embedding a learning capacity in the coach can have an important consequence for the genetic graph itself. It can eliminate the requirement that the AICAI tutor have a complete graph to teach. The graph can be incomplete but grown by the embedded learning program when needed.

7.6 The genetic graph is not a complete theory of learning.

While the issues raised in this section are provocative, the genetic graph is by no means a complete theory of learning. Hard questions remain to be studied: When should a learning strategy be applied? How are profitable analogies, generalizations and refinements detected? What are the criteria for forgetting? Furthermore, there is an enormous amount of experimental exploration that must be done. But I believe it is clear that AICAI programs will gain increased leverage by embodying an explicit theory of the learner.

8. Conclusions

My interest in the evolution of a learner's knowledge was inspired by Piaget, who often speaks of himself as a *genetic epistemologist*. He characterizes the fundamental problem of genetic epistemology as: "*the explanation of the construction of novelties in the development of knowledge.*" This paper has explored the construction of new knowledge in terms of a genetic graph. As a test of the effectiveness of this theory, I have described a design by which the graph can improve the tutoring and modelling of AICAI systems. I have also described a complementary design for a set of computer-based learning programs, in which the genetic processes form a separate expert operating on the learner's genetic graph.

Our next step will be to complete the implementation of an AICAI tutor based on the genetic graph approach, and experiment with the resulting system. I have little doubt that the genetic

graph will increase the effectiveness of this tutor over a comparable Expert-based system. More interesting will be the fine-grained analysis of learning that such a system makes possible. We will employ it to explore such Piagetian questions as:

- * Are there "stages" in the acquisition of these genetic processes as evidenced by certain explanation strategies proving unuseable for populations of different age and background?
- * Does tutoring "procedural assimilation" prove easier than tutoring "procedural accommodation", where the former is defined in terms of the acquisition of additional procedures implementing a known concept, that is intra-island rules linked by generalization, specialization, analogy and correction links; while the latter represents the acquisition of a new concept and the associated growth of a new island of rules?³⁷
- * Do islands define stable knowledge plateaus, providing a kind of "equilibration"?

While I do not know the answers to these questions, I believe this paper demonstrates that the formal study of learning and teaching required by AICAI research is a powerful methodology for studying fundamental questions in cognitive psychology and artificial intelligence.

37. For Piaget, "accomodation" for situations where the learner builds new structures to handle a task; "assimilation" involves situations where the adaptation of old structures proves sufficient. My definitions of procedural assimilation and procedural accomodation are intended to provide a loose analogy, wherein new structures correspond to new islands. I employ this analogy only to indicate that our procedural approach allows the exploration of precise definitions for the notions of local and global changes to a knowledge structure. Whether a more precise match of computational and Piagetian terminology is possible (or fruitful) remains to be seen.

9. References

- [Bro73] Brown, J.S., R. Burton and F. Zdybel, "A Model-Driven Question-Answering System for Mixed-Initiative Computer-Assisted Instruction", IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-3, No. 3, May 1973, pp. 248-257.
- [Bro77a] Brown, J.S., R. Burton and K. Larkin, "Representing and Using Procedural Bugs for Educational Purposes" Proceedings of 1977 Annual Conference, Association for Computing Machinery, Seattle, Oct. 1977, pp. 247-255.
- [Bro77b] Brown, J.S., and R. Burton, Diagnostic Models for Procedural Bugs in Basic Mathematical Skills, ICAI No. 10, Bolt, Beranek and Newman, August 1977,
- [Brow77] Brown, R., Use of Analogy to Achieve New Expertise, MIT AI Technical Report 403, April 1977.
- [Bur76] Burton, R. & J. S. Brown, "A Tutoring and Student Modelling Paradigm for Gaming Environments", in R. Coleman and P. Lorton, Jr. (Eds.), Computer Science and Education, ACM SIGCSE Bulletin, Vol. 8, No. 1, Feb. 1976, pp. 236-246.
- [Bur77] Burton, R. & J.S. Brown, Semantic Grammar: A Technique for Constructing Natural Language Interfaces to Instructional Systems, BBN Report No. 3587, ICAI Report No. 5, May 1977.
- [Car70] Carbonell, J., "AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction", IEEE Transactions on Man-Machine Systems, Vol. MMS-II, No. 4, December 1970.
- [Carr77a] Carr, B., Wusor II: A Computer Aided Instruction Program With Student Modelling Capabilities, MIT AI Memo 417 (LOGO Memo 45), May 1977.
- [Carr77b] Carr, B. & I. Goldstein, Overlays: A Theory of Modelling for Computer Aided Instruction, MIT AI Memo 406 (LOGO Memo 40), February 1977.
- [Col75] Collins, A., E. Warnock and J. Passafiume, "Analysis and Synthesis of Tutorial Dialogues", in G. Bower (Ed.), The Psychology of Learning and Motivation, Vol. 9, New York: Academic Press, 1975.
- [Dek76] De Kleer, Johan, Local methods for Localizing Faults in Electronic Circuits, MIT AI Memo 394, Nov., 1976.
- [Eva68] Evans, T., "A Program for the Solution of Geometric-Analogy Intelligence Test Questions", in M. Minsky (Ed.) Semantic Information Processing, Cambridge, MA: The MIT Press, 1968, pp. 271-353.
- [Gol75] Goldstein, I., "Summary of MYCROFT: A System for Understanding Simple Picture Programs", Artificial Intelligence Journal, Vol. 6, No. 3, Fall 1975.

[Gol77a] Goldstein, I. and B. Carr, "The Computer as Coach: An Athletic Paradigm for Intellectual Education", Proceedings of 1977 Annual Conference, Association for Computing Machinery, Seattle, October 1977, pp. 227-233.

[Gol77b] Goldstein, I. and E. Grimson, Annotated Production Systems: A Model for Skill Acquisition, MIT AI Memo 407 (LOGO Memo 44), February 1977.

[Kof75] Koffman, E. and S. Blount, "Artificial Intelligence and Automatic Programming in CAI", Artificial Intelligence, Vol. 6, 1975, pp. 215-234.

[Moor73] Moore, J. and A. Newell, "How Can MERLIN Understand?", in L. Gregg (Ed.), Knowledge and Cognition, Potomac, MD: Lawrence Erlbaum Associates, 1973.

[Nor76] Norman, D., Studies of Learning and Self-Contained Educational Systems, 1973-1976, University of California at San Diego, Center for Human Information Processing, Report No. 7601, March 1976.

[Nor76] Norman, D., D. Gentner and A. Stevens, "Comments on Learning: Schemata and Memory Representation", in D. Klahr (Ed.), Cognition and Instruction, Hillsdale, NJ: Erlbaum Associates, 1976.

[Sac75] Sacerdoti, E., "The Nonlinear Nature of Plans", Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, Georgia, USSR, 1975, pp. 206-218.

[Sel74] Self, J., "Student Models in Computer-Aided Instruction", International Journal of Man-Machine Studies, Vol. 6, 1974, pp. 261-276.

[Sel77] Self, J., "Concept Teaching". Artificial Intelligence Journal, Vol. 9, No. 2, Oct. 1977, pp. 197-221.

[Sle75] Sleeman, D., "A Problem-Solving Monitor for a Deductive Reasoning Task", International Journal of Man-Machine Studies, Vol. 7, 1975, pp. 183-211.

[Sle77] Sleeman, D., "A System Which Allows Students to Explore Algorithms", Proceedings of the Fifth International Joint Conference on Artificial Intelligence, August 1977, pp. 780-786.

[Smi75] Smith, R.L., H. Graves, L.H. Blaine, & V.G. Marinov, "Computer-Assisted Axiomatic Mathematics: Informal Rigor", in O. Lecareme & R. Lewis (Eds.), Computers in Education Part I: IFIP, Amsterdam: North Holland, 1975.

[Sta76] Stansfield, J., B. Carr and I. Goldstein, Wumpus Advisor I: A First Implementation of a Program that Tutors Logical and Probabilistic Reasoning Skills, MIT AI Laboratory Memo No. 381, Sept. 1976.

[Ste77] Stevens, A. and A. Collins, "The Goal Structure of a Socratic Tutor", Proceedings of 1977 Annual Conference, Association for Computing Machinery, Seattle, October 1977, pp. 256-263.

[Sus75] Sussman, G., A Computational Model of Skill Acquisition, New York: American Elsevier, 1975.

[Wes77] Wescourt, K., M. Beard and L. Gould, "Knowledge-Based Adaptive Curriculum Sequencing for CAI: Application of a Network Representation", Proceedings of 1977 Annual Conference, Association for Computing Machinery, October 1977, pp. 234-240.

[Win75] Winston, P., "Learning Structural Descriptions from Examples", in P. Winston (Ed.), The Psychology of Computer Vision, New York: McGraw-Hill, 1975, pp. 157-209.

[Yob75] Yob, G., "Hunt the Wumpus", Creative Computing, September/October, 1975, pp. 51-54.

